arXiv:1403.8144v1 [cs.LG] 31 Mar 2014

Coding for Random Projections and Approximate Near Neighbor Search

Ping Li
Department of Statistics & Biostatistics
Department of Computer Science
Rutgers University
Piscataway, NJ 08854, USA
pingli@stat.rutgers.edu

Michael Mitzenmacher
School of Engineering and Applied Sciences
Harvard University
Cambridge, MA 02138, USA
michaelm@eecs.harvard.edu

Anshumali Shrivastava
Department of Computer Science
Cornell University
Ithaca, NY 14853, USA
anshu@cs.cornell.edu

Abstract

This technical note compares two coding (quantization) schemes for random projections in the context of sub-linear time approximate near neighbor search. The first scheme is based on uniform quantization [4] while the second scheme utilizes a uniform quantization plus a uniformly random offset [1] (which has been popular in practice). The prior work [4] compared the two schemes in the context of similarity estimation and training linear classifiers, with the conclusion that the step of random offset is not necessary and may hurt the performance (depending on the similarity level). The task of near neighbor search is related to similarity estimation with importance distinctions and requires own study. In this paper, we demonstrate that in the context of near neighbor search, the step of random offset is not needed either and may hurt the performance (sometimes significantly so, depending on the similarity and other parameters).

For approximate near neighbor search, when the target similarity level is high (e.g., correlation > 0.85), our analysis suggest to use a uniform quantization to build hash tables, with a bin width $w = 1 \sim 1.5$. On the other hand, when the target similarity level is not that high, it is preferable to use larger w values (e.g., $w \geq 2 \sim 3$). This is equivalent to say that it suffices to use only a small number of bits (or even just 1 bit) to code each hashed value in the context of sublinear time near neighbor search. An extensive experimental study on two reasonably large datasets confirms the theoretical finding.

Coding for building hash tables is a different task from coding for similarity estimation. For near neighbor search, we need coding of the projected data to determine which buckets the data points should be placed in (and the coded values are not stored). For similarity estimation, the purpose of coding is for accurately estimating the similarities using small storage space. Therefore, if necessary, we can actually code the projected data twice (with different bin widths).

In this paper, we do not study the important issue of "re-ranking" of retrieved data points by using estimated similarities. That step is needed when exact (all pairwise) similarities can not be practically stored or computed on the fly. In a concurrent work [5], we demonstrate that the retrieval accuracy can be further improved by using nonlinear estimators of the similarities based on a 2-bit coding scheme.

1 Introduction

This paper focuses on the comparison of two quantization schemes for random projections in the context of sublinear time near neighbor search. The task of near neighbor search is to identify a set of data points which are "most similar" (in some measure of similarity) to a query data point. Efficient algorithms for near neighbor search have numerous applications in search, databases, machine learning, recommending systems, computer vision, etc. Developing efficient algorithms for finding near neighbors has been an active research topic since the early days of modern computing [2]. Near neighbor search with extremely high-dimensional data (e.g., texts or images) is still a challenging task and an active research problem.

Among many types of similarity measures, the (squared) Euclidian distance (denoted by d) and the correlation (denoted by ρ) are most commonly used. Without loss of generality, consider two high-dimensional data vectors $u, v \in \mathbb{R}^D$. The squared Euclidean distance and correlation are defined as follows:

$$d = \sum_{i=1}^{D} |u_i - v_i|^2, \qquad \rho = \frac{\sum_{i=1}^{D} u_i v_i}{\sqrt{\sum_{i=1}^{D} u_i^2}\sqrt{\sum_{i=1}^{D} v_i^2}} \tag{1}$$

In practice, it appears that the correlation is more often used than the distance, partly because $|\rho|$ is nicely normalized within 0 and 1. In fact, in this study, we will assume that the marginal l_2 norms $\sum_{i=1}^{D} |u_i|^2$ and $\sum_{i=1}^{D} |v_i|^2$ are known. This is a reasonable assumption. Computing the marginal l_2 norms only requires scanning the data once, which is anyway needed during the data collection process. In machine learning practice, it is common to first normalize the data (to have unit l_2 norm) before feeding the data to classification (e.g., SVM) or clustering (e.g., K-means) algorithms.

For convenience, throughout this paper, we assume unit l_2 norms, i.e.,

$$\rho = \frac{\sum_{i=1}^{D} u_i v_i}{\sqrt{\sum_{i=1}^{D} u_i^2}\sqrt{\sum_{i=1}^{D} v_i^2}} = \sum_{i=1}^{D} u_i v_i, \qquad \text{where} \quad \sum_{i=1}^{D} u_i^2 = \sum_{i=1}^{D} v_i^2 = 1 \tag{2}$$

1.1 Random Projections

As an effective tool for dimensionality reduction, the idea of random projections is to multiply the data, e.g., $u, v \in \mathbb{R}^D$, with a random normal projection matrix $\mathbf{R} \in \mathbb{R}^{D \times k}$ (where $k \ll D$), to generate:

$$x = u \times \mathbf{R} \in \mathbb{R}^k, \quad y = v \times \mathbf{R} \in \mathbb{R}^k, \quad \mathbf{R} = \{r_{ij}\}_{i=1 j=1}^{D \quad k}, \quad r_{ij} \sim N(0,1) \text{ i.i.d.} \tag{3}$$

The method of random projections has become popular for large-scale machine learning applications such as classification, regression, matrix factorization, singular value decomposition, near neighbor search, etc.

The potential benefits of coding with a small number of bits arise because the (uncoded) projected data, $x_j = \sum_{i=1}^{D} u_i r_{ij}$ and $y_j = \sum_{i=1}^{D} v_i r_{ij}$, being real-valued numbers, are neither convenient/economical for storage and transmission, nor well-suited for indexing. The focus of this paper is on approximate (sublinear time) near neighbor search in the framework of *locality sensitive hashing* [3]. In particular, we will compare two **coding** (quantization) schemes of random projections [1, 4] in the context of near neighbor search.

1.2 Uniform Quantization

The recent work [4] proposed an intuitive coding scheme, based on a simple uniform quantization:

$$h_w^{(j)}(u) = \lfloor x_j / w \rfloor, \qquad h_w^{(j)}(v) = \lfloor y_j / w \rfloor \tag{4}$$

where $w > 0$ is the bin width and $\lfloor . \rfloor$ is the standard floor operation.

The following theorem is proved in [4] about the collision probability $P_w = \mathbf{Pr}\left(h_w^{(j)}(u) = h_w^{(j)}(v)\right)$.

Theorem 1

$$P_w = \mathbf{Pr}\left(h_w^{(j)}(u) = h_w^{(j)}(v)\right) = 2\sum_{i=0}^{\infty}\int_{iw}^{(i+1)w}\phi(z)\left[\Phi\left(\frac{(i+1)w-\rho z}{\sqrt{1-\rho^2}}\right) - \Phi\left(\frac{iw-\rho z}{\sqrt{1-\rho^2}}\right)\right]dz \quad (5)$$

In addition, P_w is a monotonically increasing function of ρ.

The fact that P_w is a monotonically increasing function of ρ makes (4) a suitable coding scheme for approximate near neighbor search in the general framework of locality sensitive hashing (LSH).

1.3 Uniform Quantization with Random Offset

[1] proposed the following well-known coding scheme, which uses windows and a random offset:

$$h_{w,q}^{(j)}(u) = \left\lfloor \frac{x_j + q_j}{w} \right\rfloor, \qquad h_{w,q}^{(j)}(v) = \left\lfloor \frac{y_j + q_j}{w} \right\rfloor \quad (6)$$

where $q_j \sim uniform(0, w)$. [1] showed that the collision probability can be written as

$$P_{w,q} = \mathbf{Pr}\left(h_{w,q}^{(j)}(u) = h_{w,q}^{(j)}(v)\right) = \int_0^w \frac{1}{\sqrt{d}} 2\phi\left(\frac{t}{\sqrt{d}}\right)\left(1 - \frac{t}{w}\right)dt \quad (7)$$

where $d = ||u-v||^2 = 2(1-\rho)$ is the Euclidean distance between u and v. Compared with (6), the scheme (4) does not use the additional randomization with $q \sim uniform(0, w)$ (i.e., the offset). [4] elaborated the following advantages of (4) in the context of similarity estimation:

1. Operationally, h_w is simpler than $h_{w,q}$.
2. With a fixed w, h_w is always more accurate than $h_{w,q}$, often significantly so.
3. For each coding scheme, one can separately find the optimum bin width w. The optimized h_w is also more accurate than optimized $h_{w,q}$, often significantly so.
4. h_w requires a smaller number of bits than $h_{w,q}$.

In this paper, we will compare $h_{w,q}$ with h_w in the context of sublinear time near neighbor search.

1.4 Sublinear Time c-Approximate Near Neighbor Search

Consider a data vector u. Suppose there exists another vector whose Euclidian distance ($\sqrt{d}$) from u is at most $\sqrt{d_0}$ (the target distance). The goal of *c-approximate $\sqrt{d_0}$-near neighbor* algorithms is to return data vectors (with high probability) whose Euclidian distances from u are at most $c \times \sqrt{d_0}$ with $c > 1$.

Recall that, in our definition, $d = 2(1-\rho)$ is the squared Euclidian distance. To be consistent with [1], we present the results in terms of $\sqrt{d}$. Corresponding to the target distance $\sqrt{d_0}$, the target similarity ρ_0 can be computed from $d_0 = 2(1-\rho_0)$ i.e., $\rho_0 = 1 - d_0/2$. To simplify the presentation, we focus on $\rho \geq 0$ (as is common in practice), i.e., $0 \leq d \leq 2$. Once we fix a target similarity ρ_0, c can not exceed a certain value:

$$c\sqrt{2(1-\rho_0)} \leq \sqrt{2} \Longrightarrow c \leq \sqrt{\frac{1}{1-\rho_0}} \quad (8)$$

For example, when $\rho_0 = 0.5$, we must have $1 \leq c \leq \sqrt{2}$.

Under the general framework, the performance of an LSH algorithm largely depends on the difference (gap) between the two collision probabilities $P^{(1)}$ and $P^{(2)}$ (respectively corresponding to $\sqrt{d_0}$ and $c\sqrt{d_0}$):

$$P_w^{(1)} = \mathbf{Pr}\left(h_w(u) = h_w(v)\right) \qquad \text{when } d = ||u-v||_2^2 = d_0 \tag{9}$$

$$P_w^{(2)} = \mathbf{Pr}\left(h_w(u) = h_w(v)\right) \qquad \text{when } d = ||u-v||_2^2 = c^2 d_0 \tag{10}$$

Corresponding to $h_{w,q}$, the collision probabilities $P_{w,q}^{(1)}$ and $P_{w,q}^{(2)}$ are analogously defined.

A larger difference between $P^{(1)}$ and $P^{(2)}$ implies a more efficient LSH algorithm. The following "G" values (G_w for h_w and $G_{w,q}$ for $h_{w,q}$) characterize the gaps:

$$G_w = \frac{\log 1/P_w^{(1)}}{\log 1/P_w^{(2)}}, \qquad G_{w,q} = \frac{\log 1/P_{w,q}^{(1)}}{\log 1/P_{w,q}^{(2)}} \tag{11}$$

A smaller G (i.e., larger difference between $P^{(1)}$ and $P^{(2)}$) leads to a potentially more efficient LSH algorithm and $\rho < \frac{1}{c}$ is particularly desirable [3]. The general theory says the query time for c-approximate d_0-near neighbor is dominated by $O(N^G)$ distance evaluations, where N is the total number of data vectors in the collection. This is better than $O(N)$, the cost of a linear scan.

2 Comparison of the Collision Probabilities

To help understand the intuition why h_w may lead to better performance than $h_{w,q}$, in this section we examine their collision probabilities P_w and $P_{w,q}$, which can be expressed in terms of the standard normal pdf and cdf functions: $\phi(x) = \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}}$ and $\Phi(x) = \int_{-\infty}^{x} \phi(x)dx$,

$$P_{w,q} = \mathbf{Pr}\left(h_{w,q}^{(j)}(u) = h_{w,q}^{(j)}(v)\right) = 2\Phi\left(\frac{w}{\sqrt{d}}\right) - 1 - \frac{2}{\sqrt{2\pi}w/\sqrt{d}} + \frac{2}{w/\sqrt{d}}\phi\left(\frac{w}{\sqrt{d}}\right) \tag{12}$$

$$P_w = \mathbf{Pr}\left(h_w^{(j)}(u) = h_w^{(j)}(v)\right) = 2\sum_{i=0}^{\infty} \int_{iw}^{(i+1)w} \phi(z)\left[\Phi\left(\frac{(i+1)w - \rho z}{\sqrt{1-\rho^2}}\right) - \Phi\left(\frac{iw - \rho z}{\sqrt{1-\rho^2}}\right)\right] dz \tag{13}$$

It is clear that $P_{w,q} \to 1$ as $w \to \infty$.

Figure 1 plots both P_w and $P_{w,q}$ for selected ρ values. The difference between P_w and $P_{w,q}$ becomes apparent when w is not small. For example, when $\rho = 0$, P_w quickly approaches the limit 0.5 while $P_{w,q}$ keeps increasing (to 1) as w increases. Intuitively, the fact that $P_{w,q} \to 1$ when $\rho = 0$, is undesirable because it means two orthogonal vectors will have the same coded value. Thus, it is not surprising that h_w will have better performance than $h_{w,q}$, for both similarity estimation and sublinear time near neighbor search.

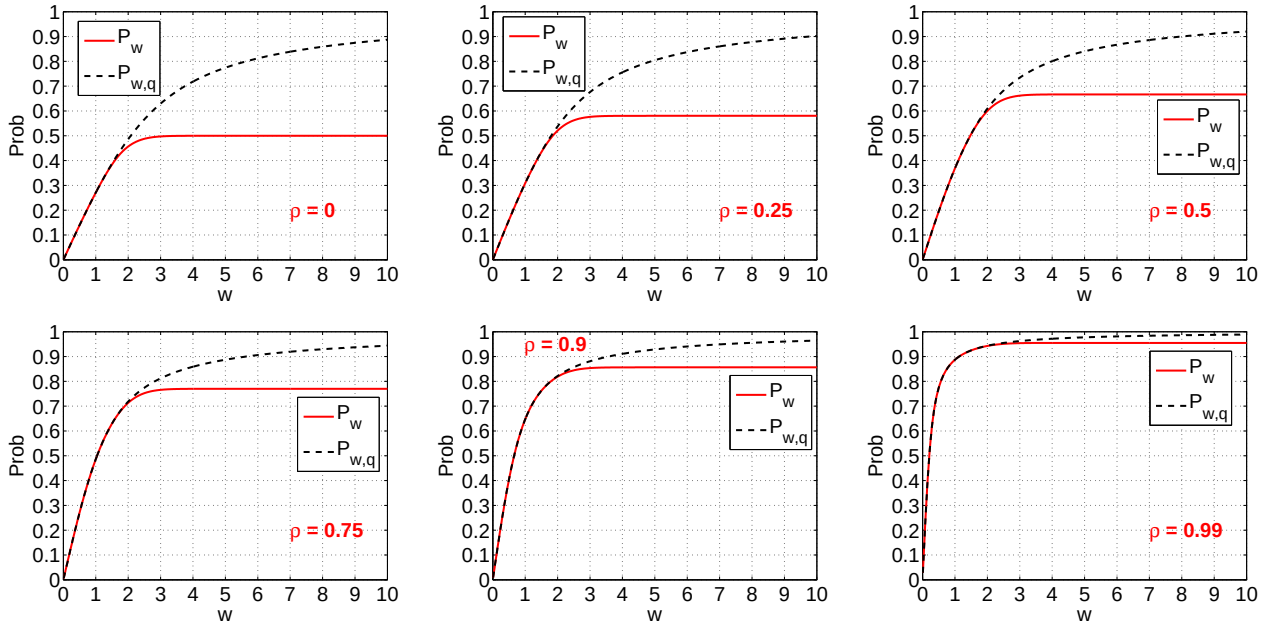

Figure 1: Collision probabilities, P_w and $P_{w,q}$, for $\rho = 0, 0.25, 0.5, 0.75, 0.9$, and 0.99. The scheme h_w has smaller collision probabilities than the scheme [1] $h_{w,q}$, especially when $w > 2$.

3 Theoretical Comparison of the Gaps

Figure 2 compares G_w with $G_{w,q}$ at their "optimum" w values, as functions of c, for a wide range of target similarity ρ_0 levels. Basically, at each c and ρ_0, we choose the w to minimize G_w and the w to minimize $G_{w,q}$. This figure illustrates that G_w is smaller than $G_{w,q}$, noticeably so in the low similarity region.

Figure 3, Figure 4, Figure 5, and Figure 6 present G_w and $G_{w,q}$ as functions of w, for $\rho_0 = 0.99$, $\rho_0 = 0.95$, $\rho_0 = 0.9$ and $\rho_0 = 0.5$, respectively. In each figure, we plot the curves for a wide range of c values. These figures illustrate where the optimum w values are obtained. Clearly, in the high similarity region, the smallest G values are obtained at low w values, especially at small c. In the low (or moderate) similarity region, the smallest G values are usually attained at relatively large w.

In practice, we normally have to pre-specify a w, for all c and ρ_0 values. In other words, the "optimum" G values presented in Figure 2 are in general not attainable. Therefore, Figure 7, Figure 8, Figure 9, and Figure 10 present G_w and $G_{w,q}$ as functions of c, for $\rho_0 = 0.99$, $\rho_0 = 0.95$, $\rho_0 = 0.9$ and $\rho_0 = 0.5$, respectively. In each figure, we plot the curves for a wide range of w values. These figures again confirm that G_w is smaller than $G_{w,q}$.

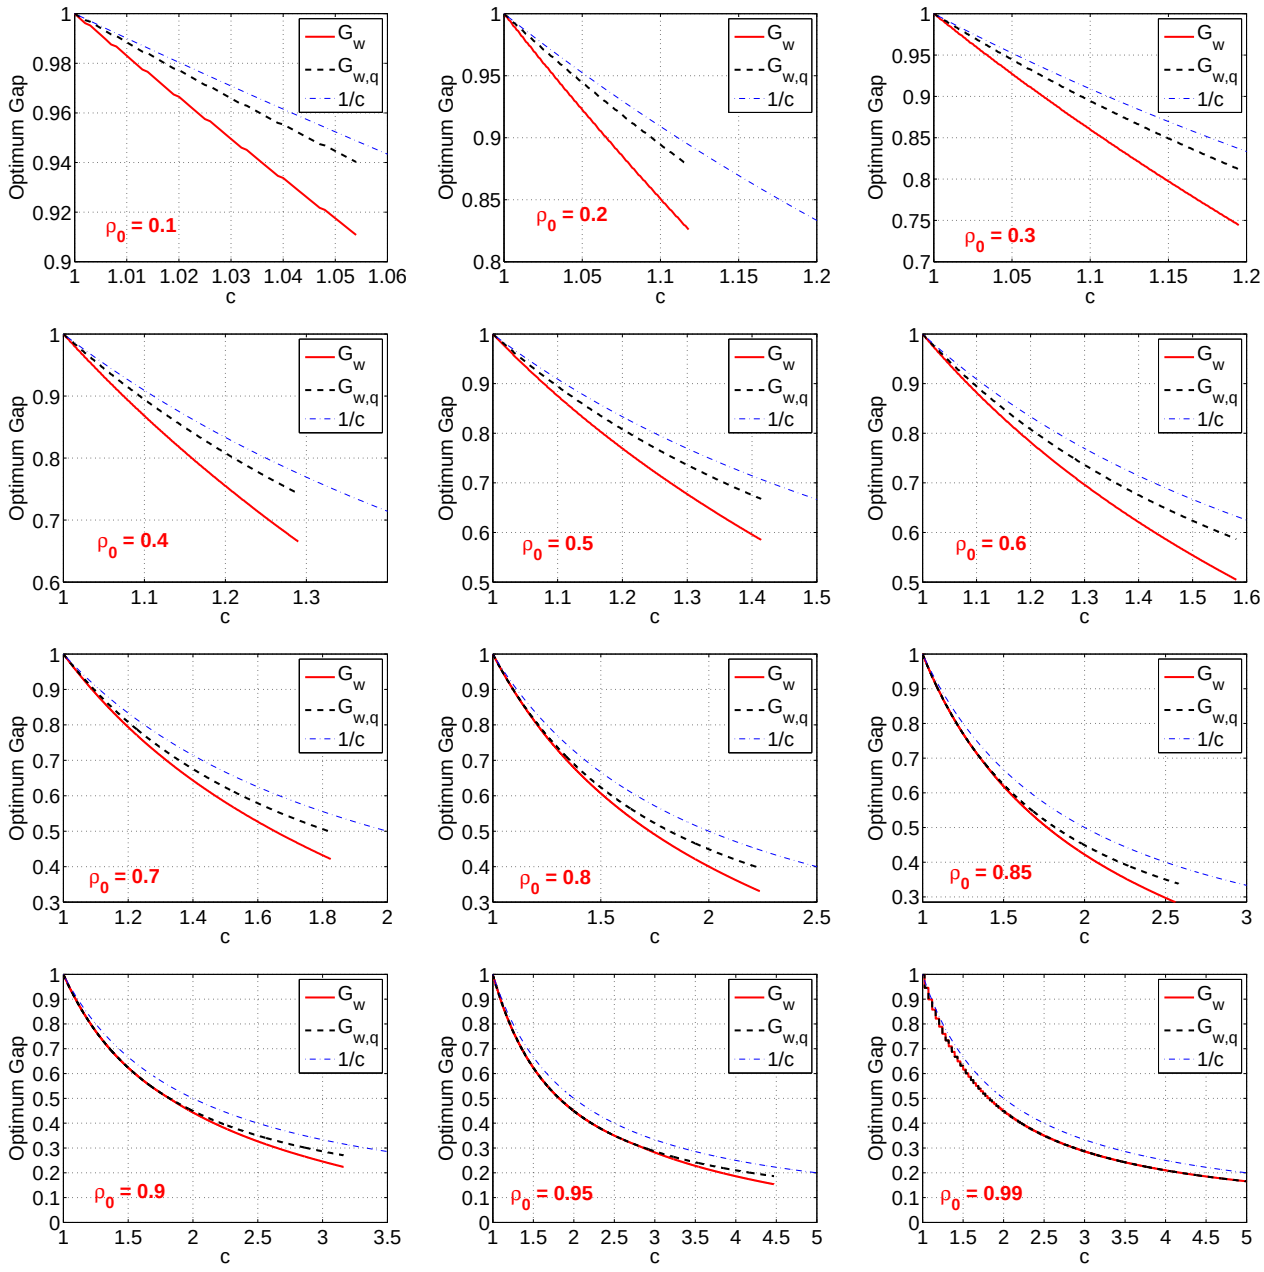

Figure 2: Comparison of the optimum gaps (smaller the better) for h_w and $h_{w,q}$. For each ρ_0 and c, we can find the smallest gaps individually for h_w and $h_{w,q}$, over the entire range of w. We can see that for all target similarity levels ρ_0, both $h_{w,q}$ and h_w exhibit better performance than $1/c$. h_w always has smaller gap than $h_{w,q}$, although in high similarity region both schemes perform similarly.

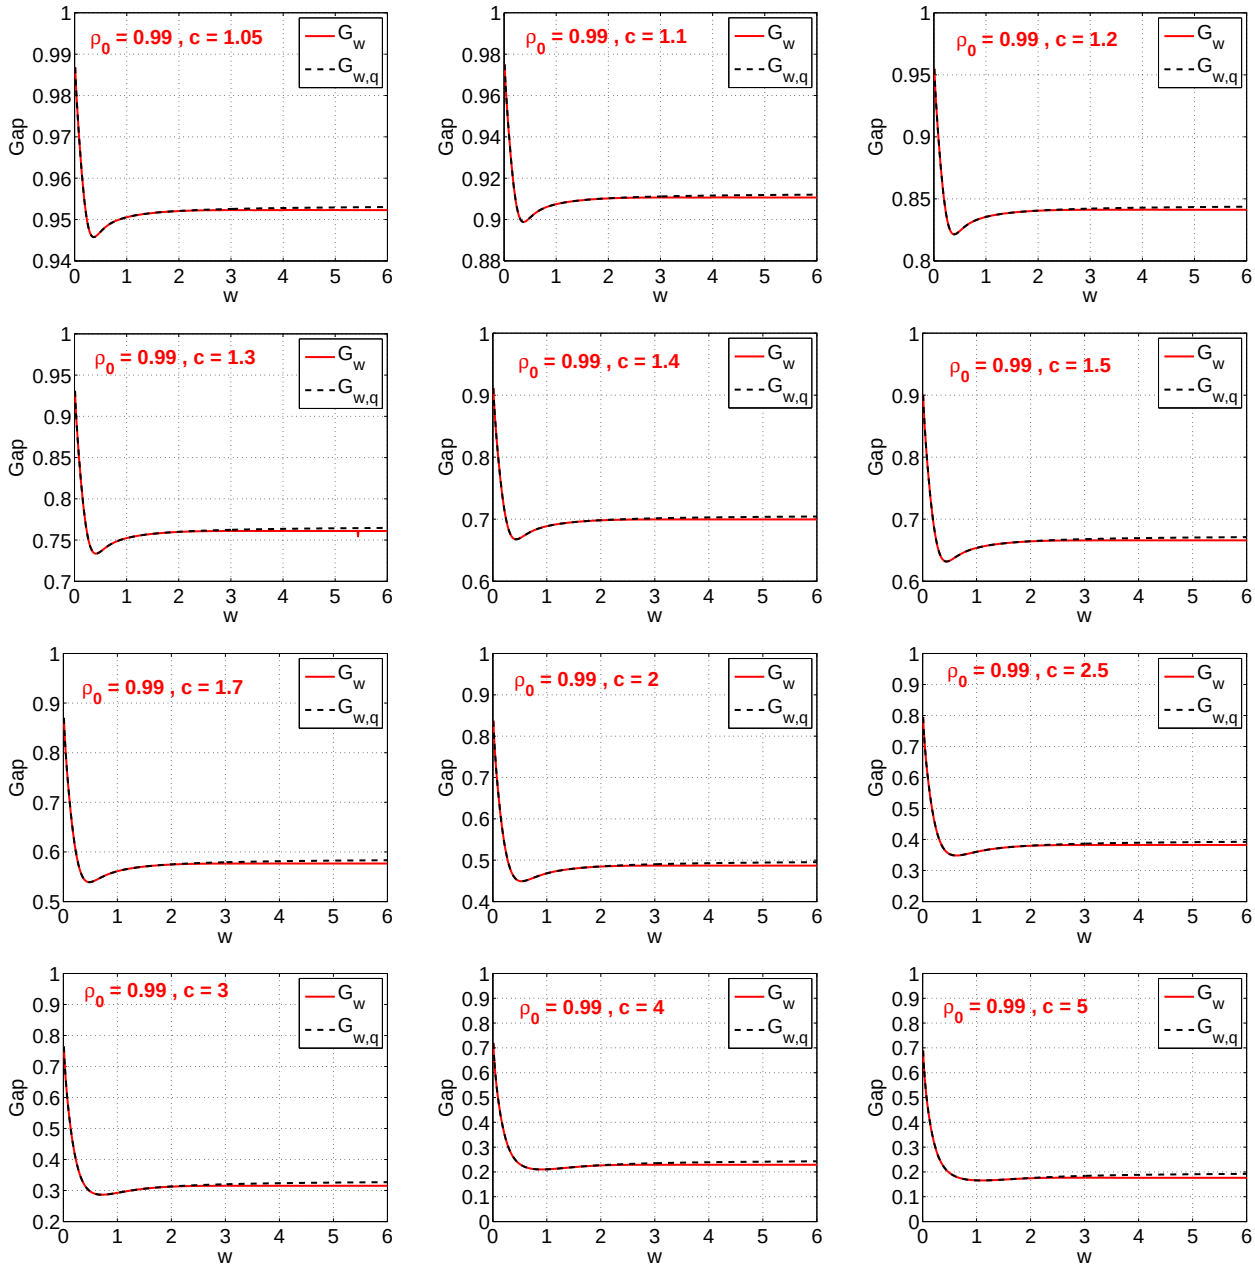

Figure 3: The gaps G_w and $G_{w,q}$ as functions of w, for $\rho_0 = 0.99$. In each panel, we plot both G_w and $G_{w,q}$ for a particular c value. The plots illustrate where the optimum w values are obtained.

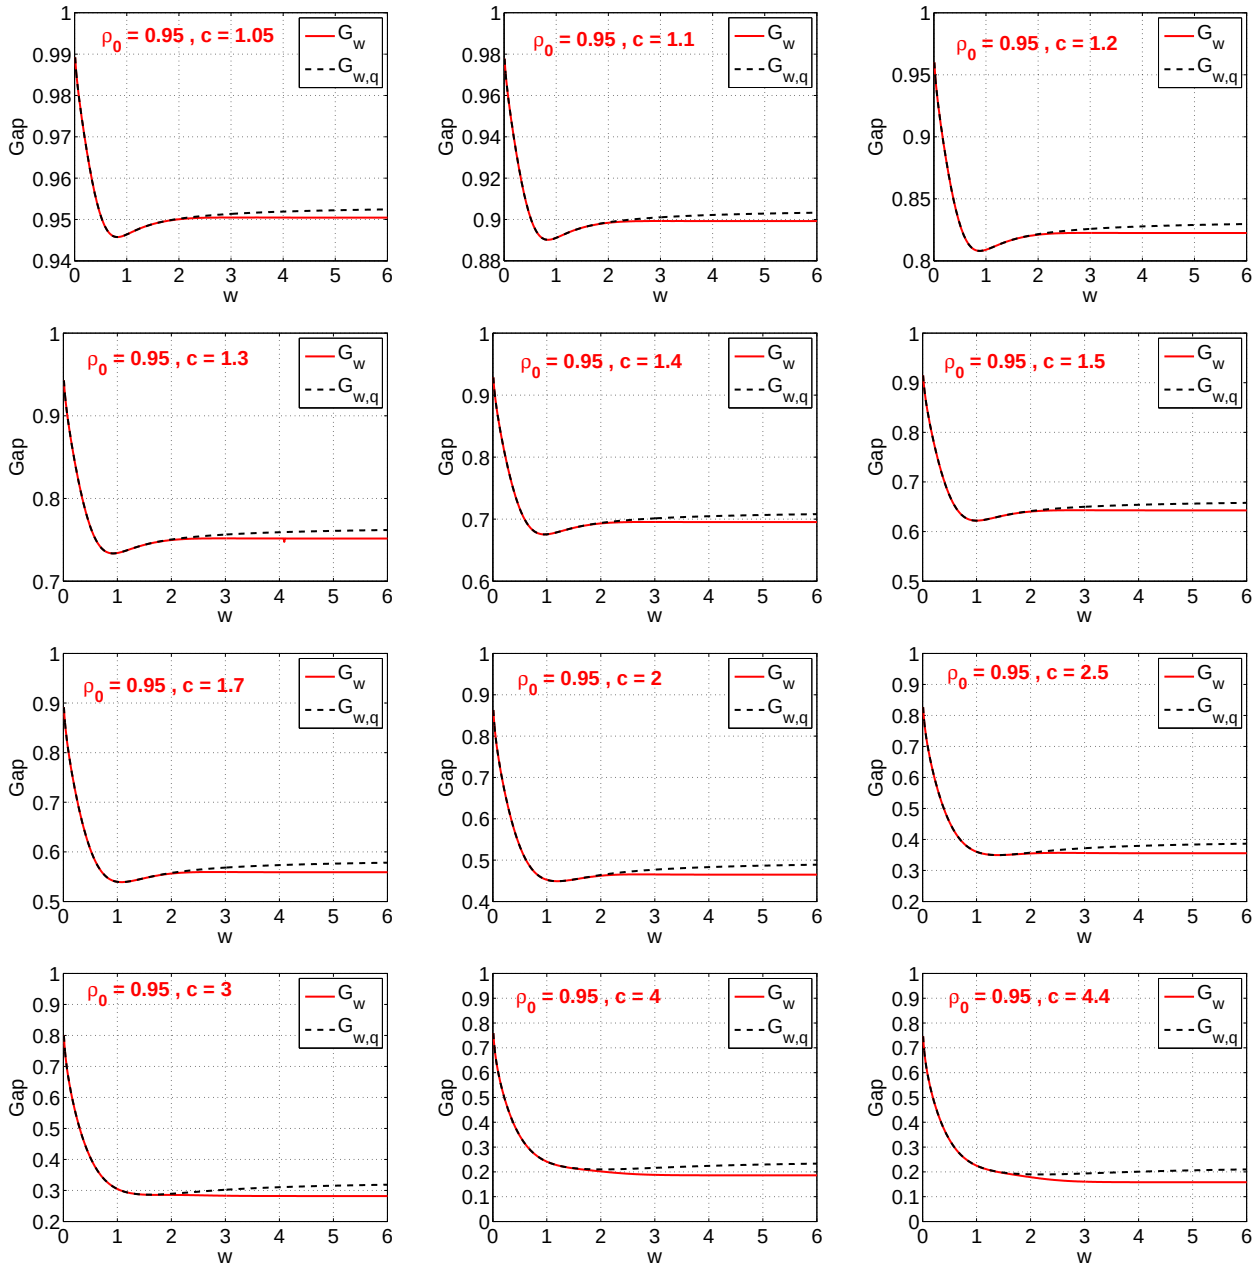

Figure 4: The gaps G_w and $G_{w,q}$ as functions of w, for $\rho_0 = 0.95$ and a range of c values.

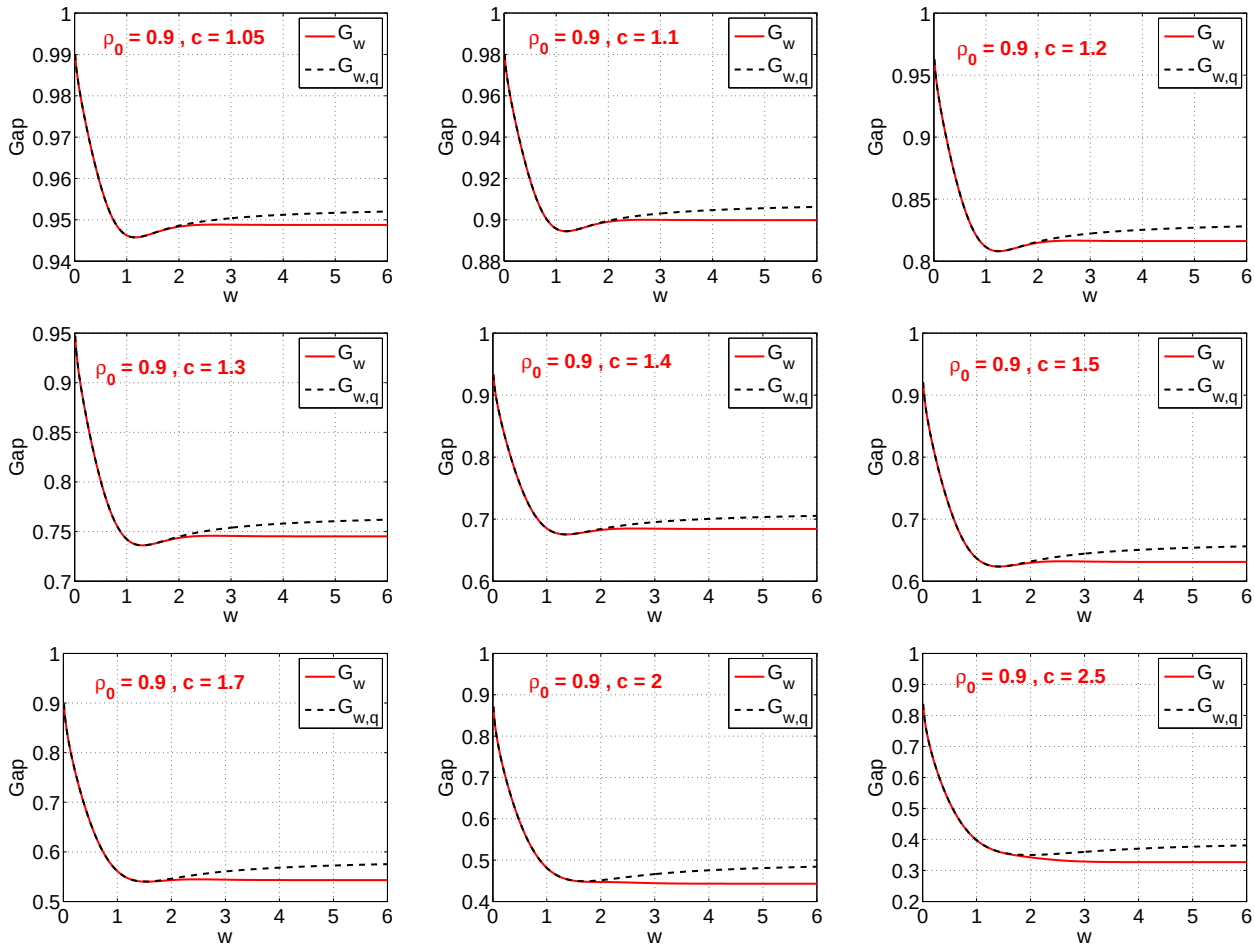

Figure 5: The gaps G_w and $G_{w,q}$ as functions of w, for $\rho_0 = 0.9$ and a range of c values.

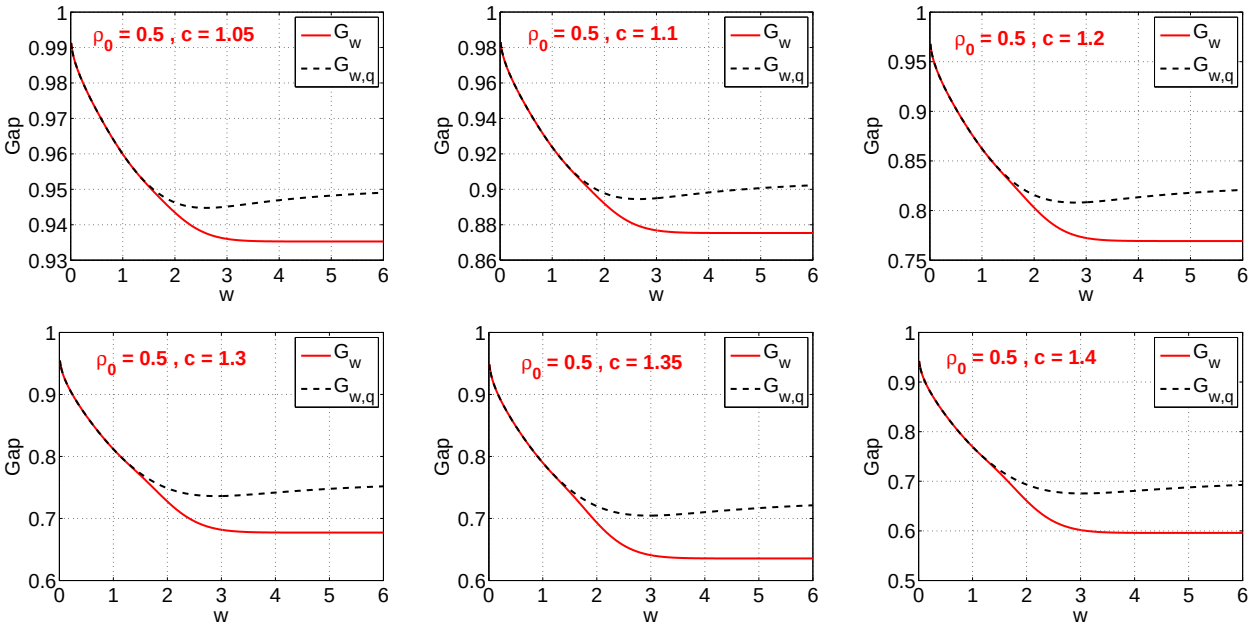

Figure 6: The gaps G_w and $G_{w,q}$ as functions of w, for $\rho_0 = 0.5$ and a range of c values.

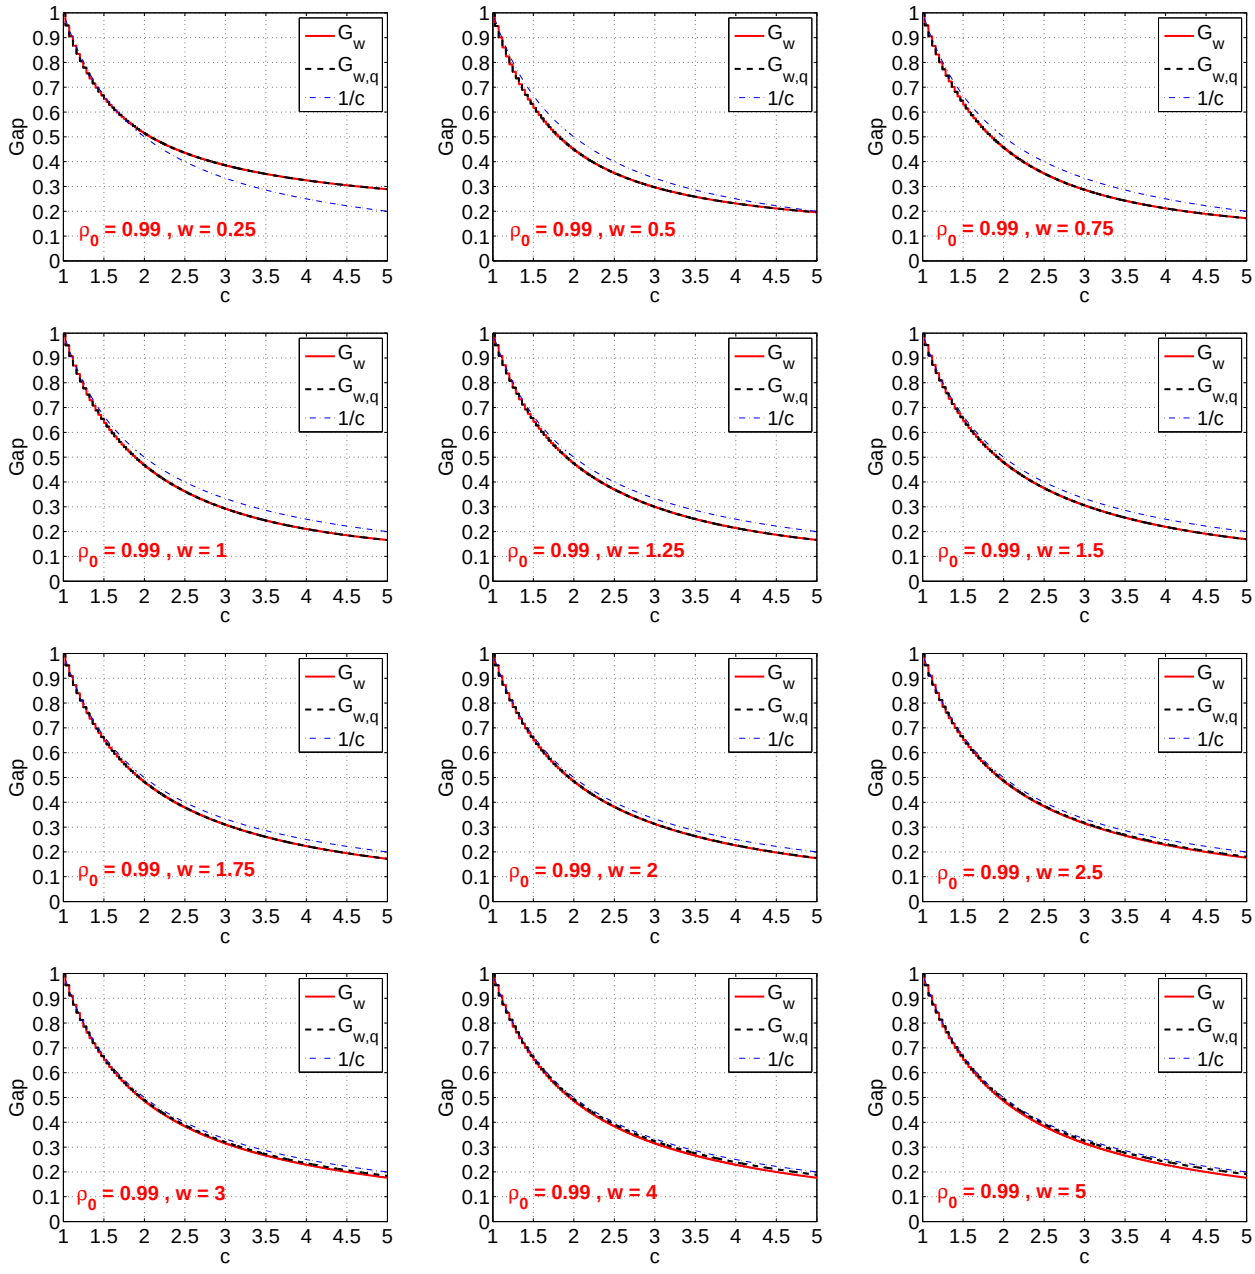

Figure 7: The gaps G_w and $G_{w,q}$ as functions of c, for $\rho_0 = 0.99$. In each panel, we plot both G_w and $G_{w,q}$ for a particular w value.

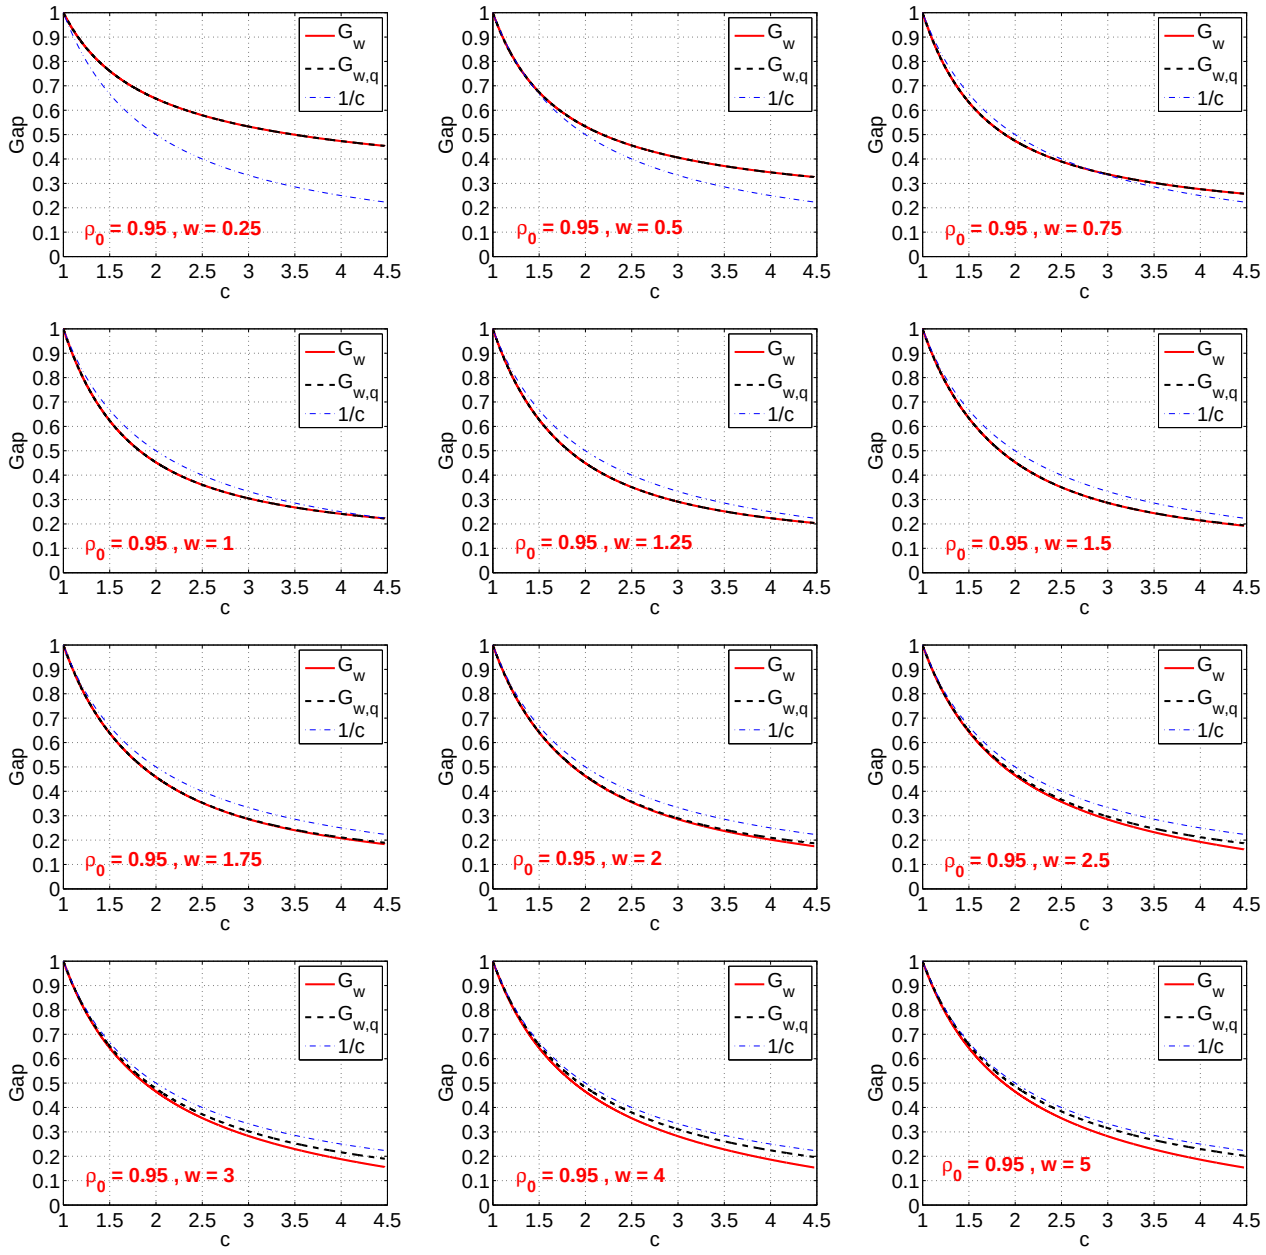

Figure 8: The gaps G_w and $G_{w,q}$ as functions of c, for $\rho_0 = 0.95$. In each panel, we plot both G_w and $G_{w,q}$ for a particular w value.

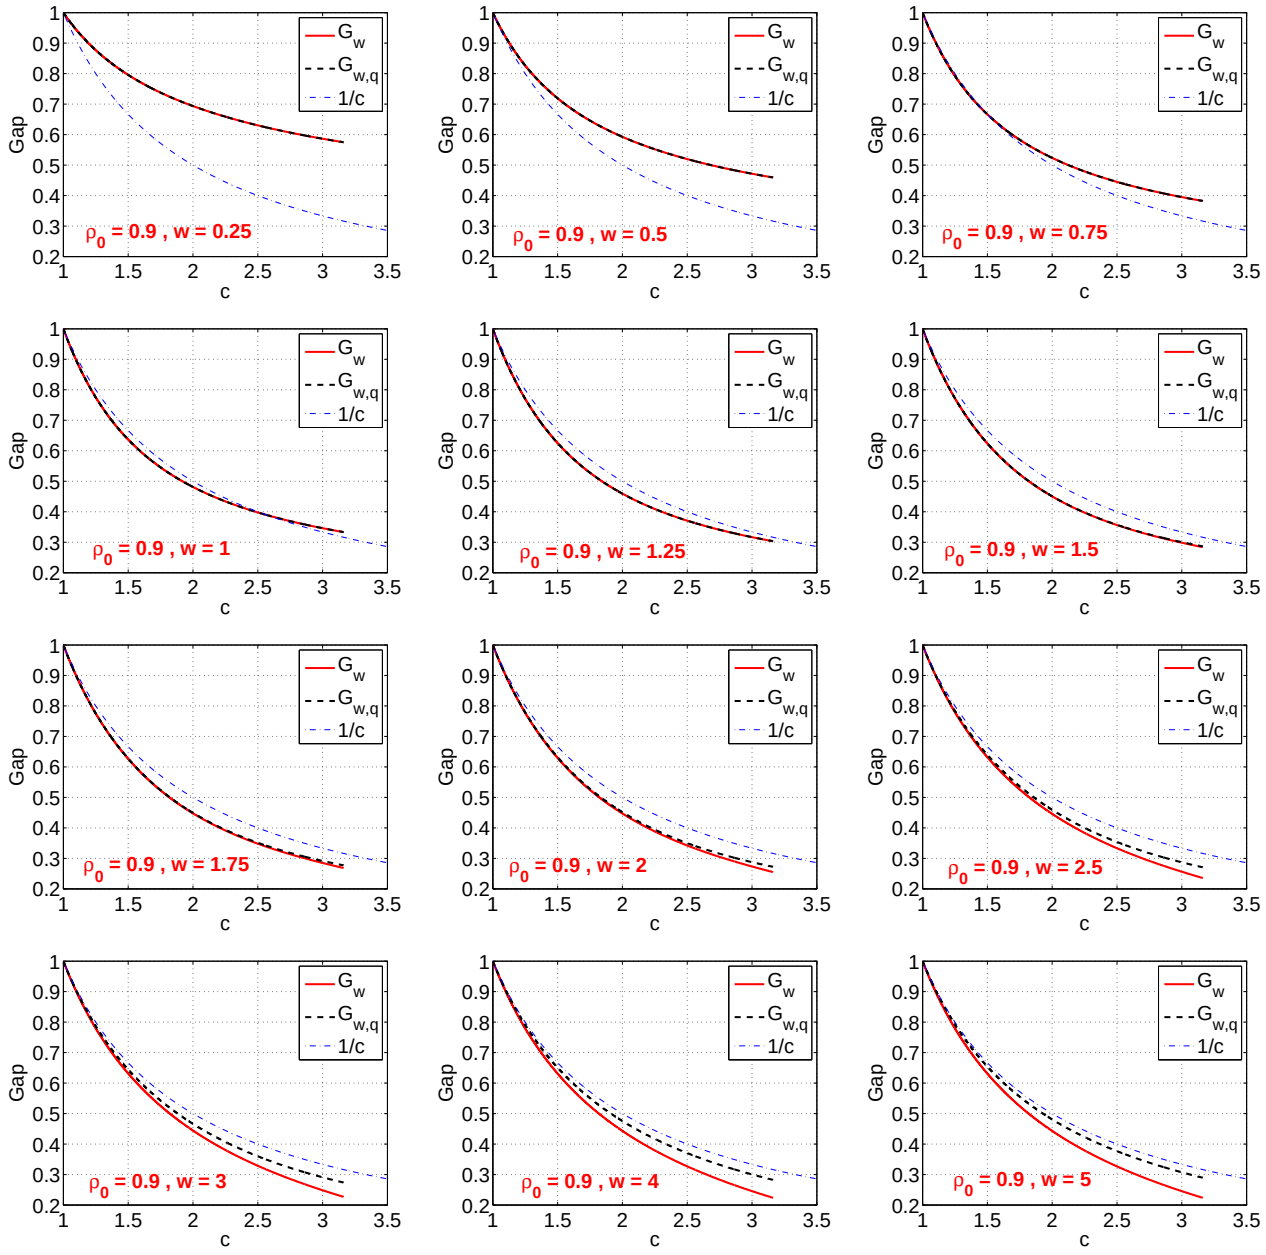

Figure 9: The gaps G_w and $G_{w,q}$ as functions of c, for $\rho_0 = 0.9$. In each panel, we plot both G_w and $G_{w,q}$ for a particular w value.

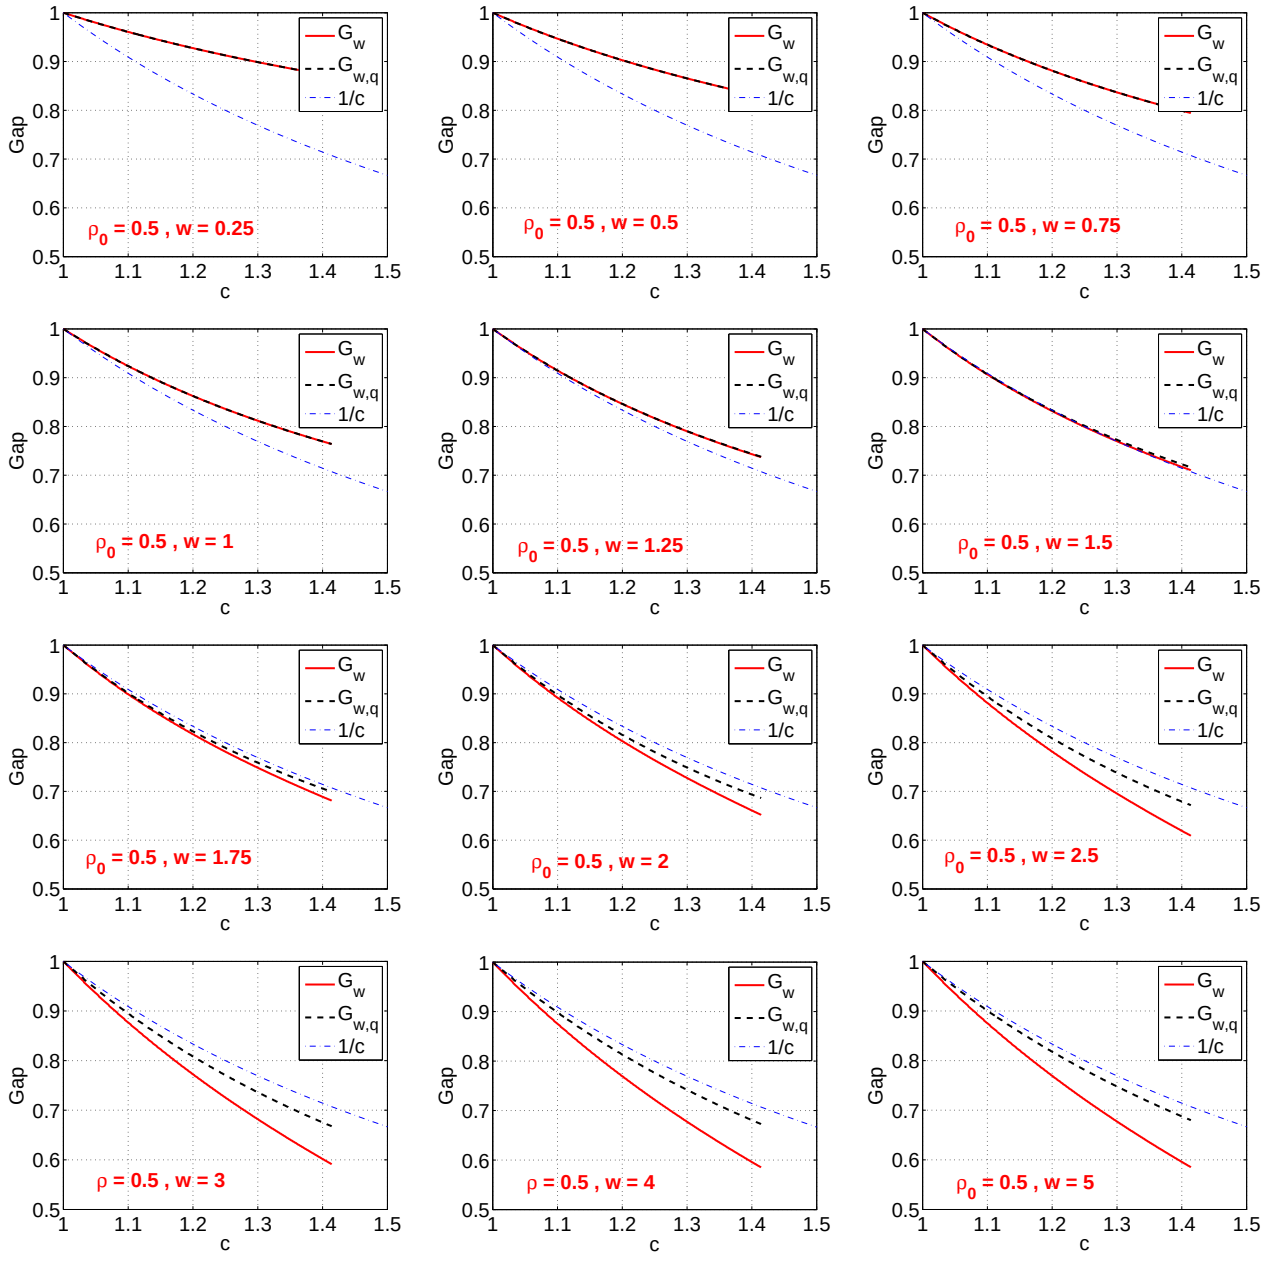

Figure 10: The gaps G_w and $G_{w,q}$ as functions of c, for $\rho_0 = 0.5$. In each panel, we plot both G_w and $G_{w,q}$ for a particular w value.

4 Optimal Gaps

To view the optimal gaps more clearly, Figure 11 and Figure 12 plot the best gaps (left panels) and the optimal w values (right panels) at which the best gaps are attained, for selected values of c and the entire range of ρ. The results can be summarized as follows

- At any ρ and c, the optimal gap $G_{w,q}$ is always at least as large as the optimal gap G_w. At relatively low similarities, the optimal $G_{w,q}$ can be substantially larger than the optimal G_w.
- When the target similarity level ρ is high (e.g., $\rho > 0.85$), for both schemes h_w and $h_{w,q}$, the optimal w values are relatively low, for example, $w = 1 \sim 1.5$ when $0.85 < \rho < 0.9$. In this region, both $h_{w,q}$ and h_w behavior similarly.
- When the target similarity level ρ is not so high, for h_w, it is best to use a large value of w, in particular $w \geq 2 \sim 3$. In comparison, for $h_{w,q}$, the optimal w values grow smoothly with decreasing ρ.

These plots again confirm the previous comparisons: (i) we should always replace $h_{w,q}$ with h_w; (ii) if we use h_w and target at very high similarity, a good choice of w might be $w = 1 \sim 1.5$; (iii) if we use h_w and the target similarity is not too high, then we can safely use $w = 2 \sim 3$.

We should also mention that, although the optimal w values for h_w appear to exhibit a "jump" in the right panels of Figure 11 and Figure 12, the choice of w does not influence the performance much, as shown in previous plots. In Figures 3 to 6, we have seen that even when the optimal w appear to approach "∞", the actual gaps are not much difference between $w = 3$ and $w \gg 3$. In the real-data evaluations in the next section, we will see the same phenomenon for h_w.

Note that the Gaussian density decays rapidly at the tail, for example, $1 - \Phi(6) = 9.9 \times 10^{-10}$. If we choose $w = 1.5$, or 2, or 3, then we just need a small number of bits to code each hashed value.

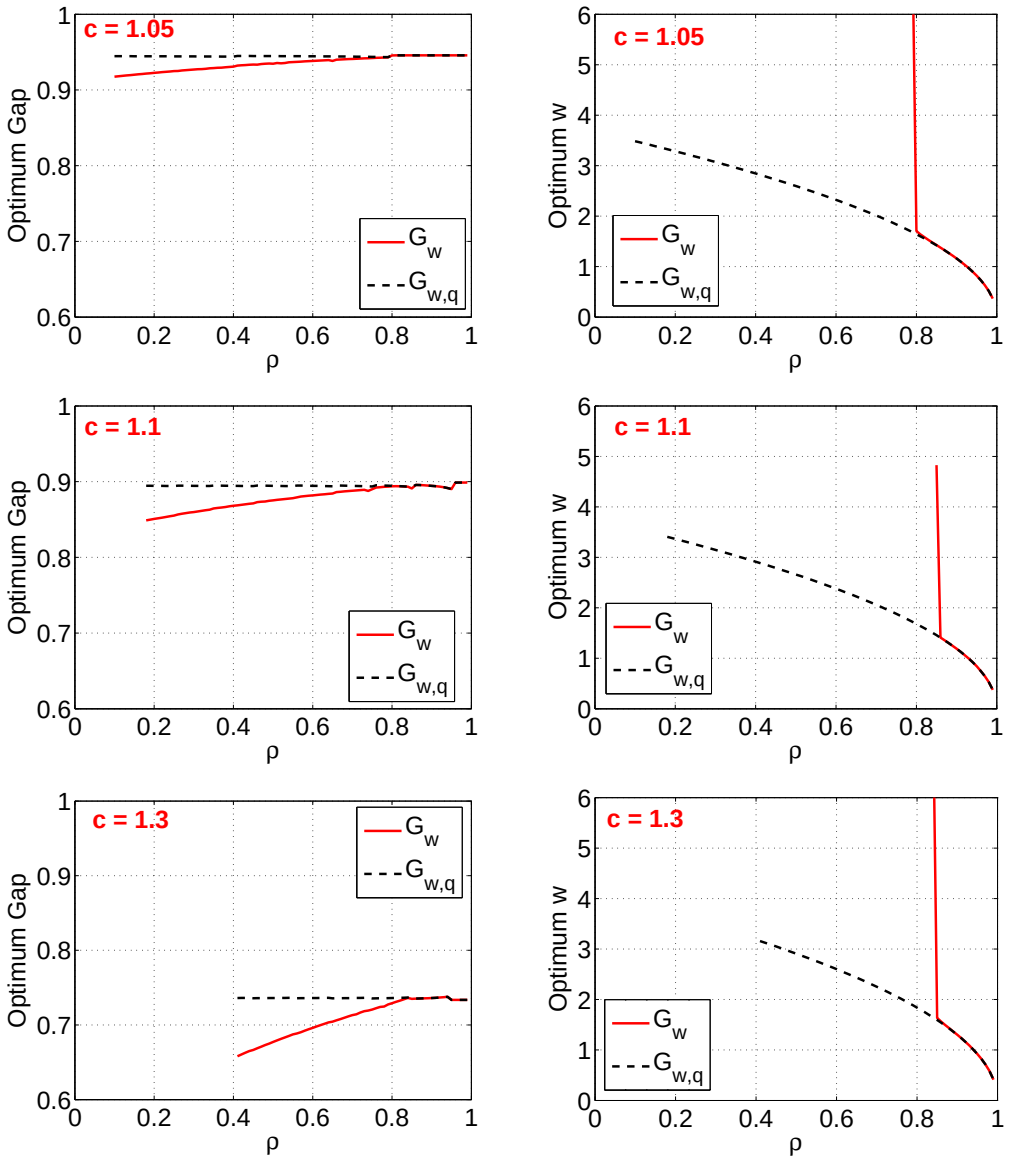

Figure 11: **Left panels**: the optimal (smallest) gaps at given c values and the entire range of ρ. We can see that $G_{w,q}$ is always larger than G_w, confirming that it is better to use h_w instead of $h_{w,q}$. **Right panels**: the optimal values of w at which the optimal gaps are attained. When the target similarity ρ is very high, it is best to use a relatively small w. When the target similarity is not that high, if we use h_w, it is best to use $w > 3$.

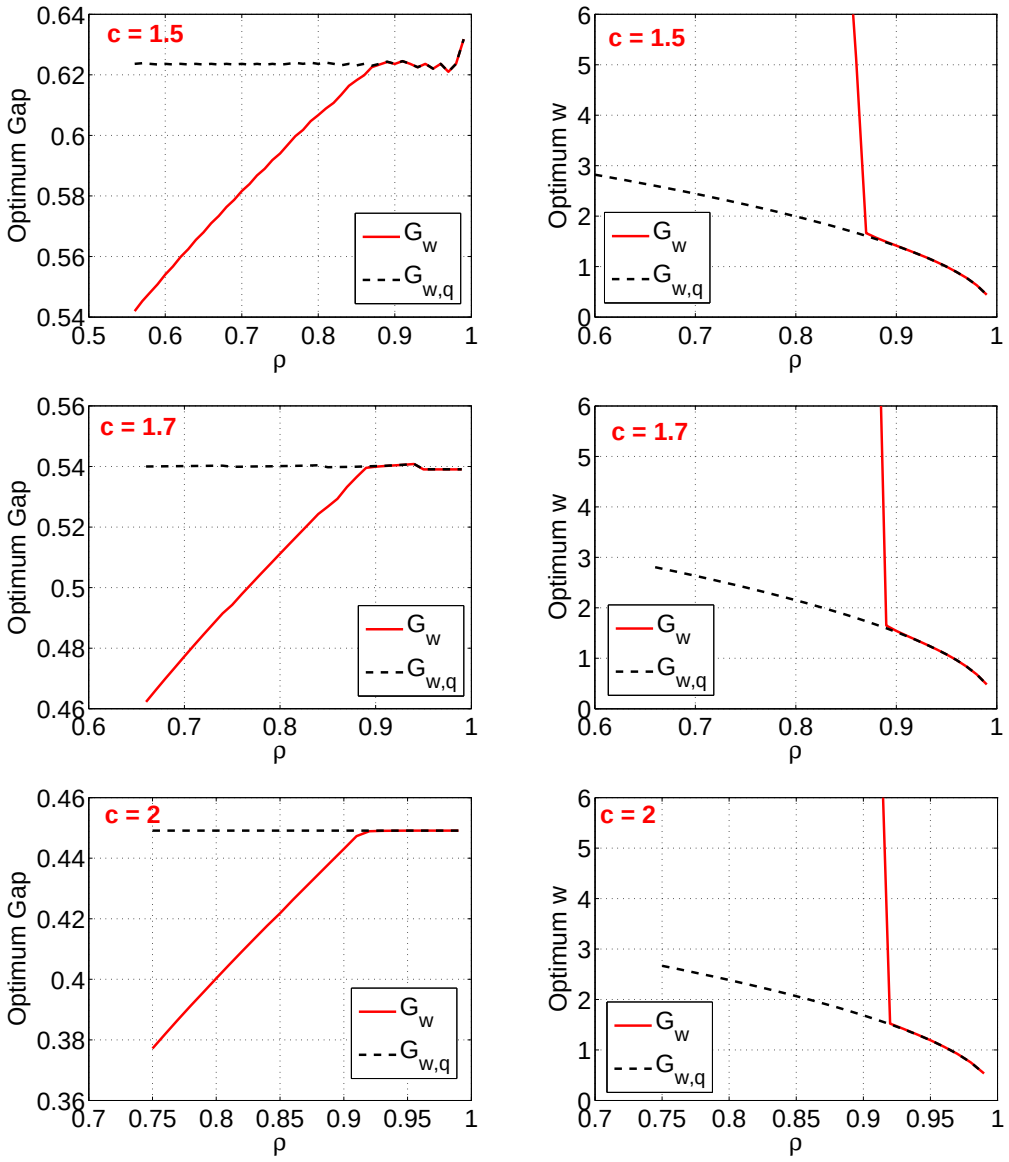

Figure 12: **Left panels**: the optimal (smallest) gaps at given c values and the entire range of ρ. We can see that $G_{w,q}$ is always larger than G_w, confirming that it is better to use h_w instead of $h_{w,q}$. **Right panels**: the optimal values of w at which the optimal gaps are attained. When the target similarity ρ is very high, it is best to use a relatively small w. When the target similarity is not that high, if we use h_w, it is best to use $w > 3$.

5 An Experimental Study

Two datasets, *Peekaboom* and *Youtube*, are used in our experiments for validating the theoretical results. *Peekaboom* is a standard image retrieval dataset, which is divided into two subsets, one with 1998 data points and another with 55599 data points. We use the larger subset for building hash tables and the smaller subset for query data points. The reported experimental results are averaged over all query data points.

Available in the UCI repository, *Youtube* is a multi-view dataset. For simplicity, we only use the largest set of audio features. The original training set, with 97934 data points, is used for building hash tables. 5000 data points, randomly selected from the original test set, are used as query data points.

We use the standard (K, L)-LSH implementation [3]. We generate $K \times L$ independent hash functions $h^{i,j}$, $i = 1$ to K, $j = 1$ to L. For each hash table j, $j = 1$ to L, we concatenate K hash functions $< h^{1,j}, h^{2,j}, h^{3,j}, ..., h^{K,j} >$. For each data point, we compute the hash values and place them (in fact, their pointers) into the appropriate buckets of the hash table i. In the query phase, we compute the hash value of the query data points using the same hash functions to find the bucket in which the query data point belongs to and only search for near neighbor among the data points in that bucket of hash table i. We repeat the process for each hash table and the final retrieved data points are the union of the retrieved data points in all the hash tables. Ideally, the number of retrieved data points will be substantially smaller than the total number of data points. We use the term *fraction retrieved* to indicate the ratio of the number of retrieved data points over the total number of data points. A smaller value of *fraction retrieved* would be more desirable.

To thoroughly evaluate the two coding schemes, we conduct extensive experiments on the two datasets, by using many combinations of K (from 3 to 40) and L (from 1 to 200). At each choice of (K, L), we vary w from 0.5 to 5. Thus, the total number of combinations is large, and the experiments are very time-consuming.

There are many ways to evaluate the performance of an LSH scheme. We could specify a threshold of similarity and only count the retrieved data points whose (exact) similarity is above the threshold as "true positives". To avoid specifying a threshold and consider the fact that in practice people often would like to retrieve the top-T nearest neighbors, we take a simple approach by computing the *recall* based on top-T neighbors. For example, suppose the number of retrieved data points is 120, among which 70 data points belong to the top-T. Then the recall value would be $70/T = 70\%$ if $T = 100$. Ideally, we hope the recalls would be as high as possible and in the meanwhile we hope to keep the fraction retrieved as low as possible.

Figure 13 presents the results on *Youtube* for $T = 100$ and target recalls from 0.1 to 0.99. In every panel, we set a target recall threshold. At every bin width w, we find the smallest *fraction retrieved* over a wide range of LSH parameters, K and L. Note that, if the target recall is high (e.g., 0.95), we basically have to effectively lower the target threshold ρ, so that we do not have to go down the re-ranked list too far. The plots show that, for high target recalls, we need to use relatively large w (e.g., $w \geq 2 \sim 3$), and for low target recalls, we should use a relatively small w (e.g., $w = 1.5$).

Figures 14 to 18 present similar results on the *Youtube* dataset for $T = 50, 20, 10, 5, 3$. We only include plots with relatively high recalls which are often more useful in practice. Figures 19 to 24 present the results on the *Peekaboom* dataset, which are essentially very similar to the results on the *Youtube* dataset.

These plots confirm the previous theoretical analysis: (i) it is essentially always better to use h_w instead of $h_{w,q}$, i.e., the random offset is not needed; (ii) when using h_w and the target recall is high (which essentially means when the target similarity is low), it is better to use a relatively large w (e.g., $w = 2 \sim 3$); (iii) when using h_w and the target recall is low, it is better to use a smaller w (e.g., $w = 1.5$); (iv) when using h_w, the influence is w is not that much as long as it is in a reasonable range, which is important in practice.

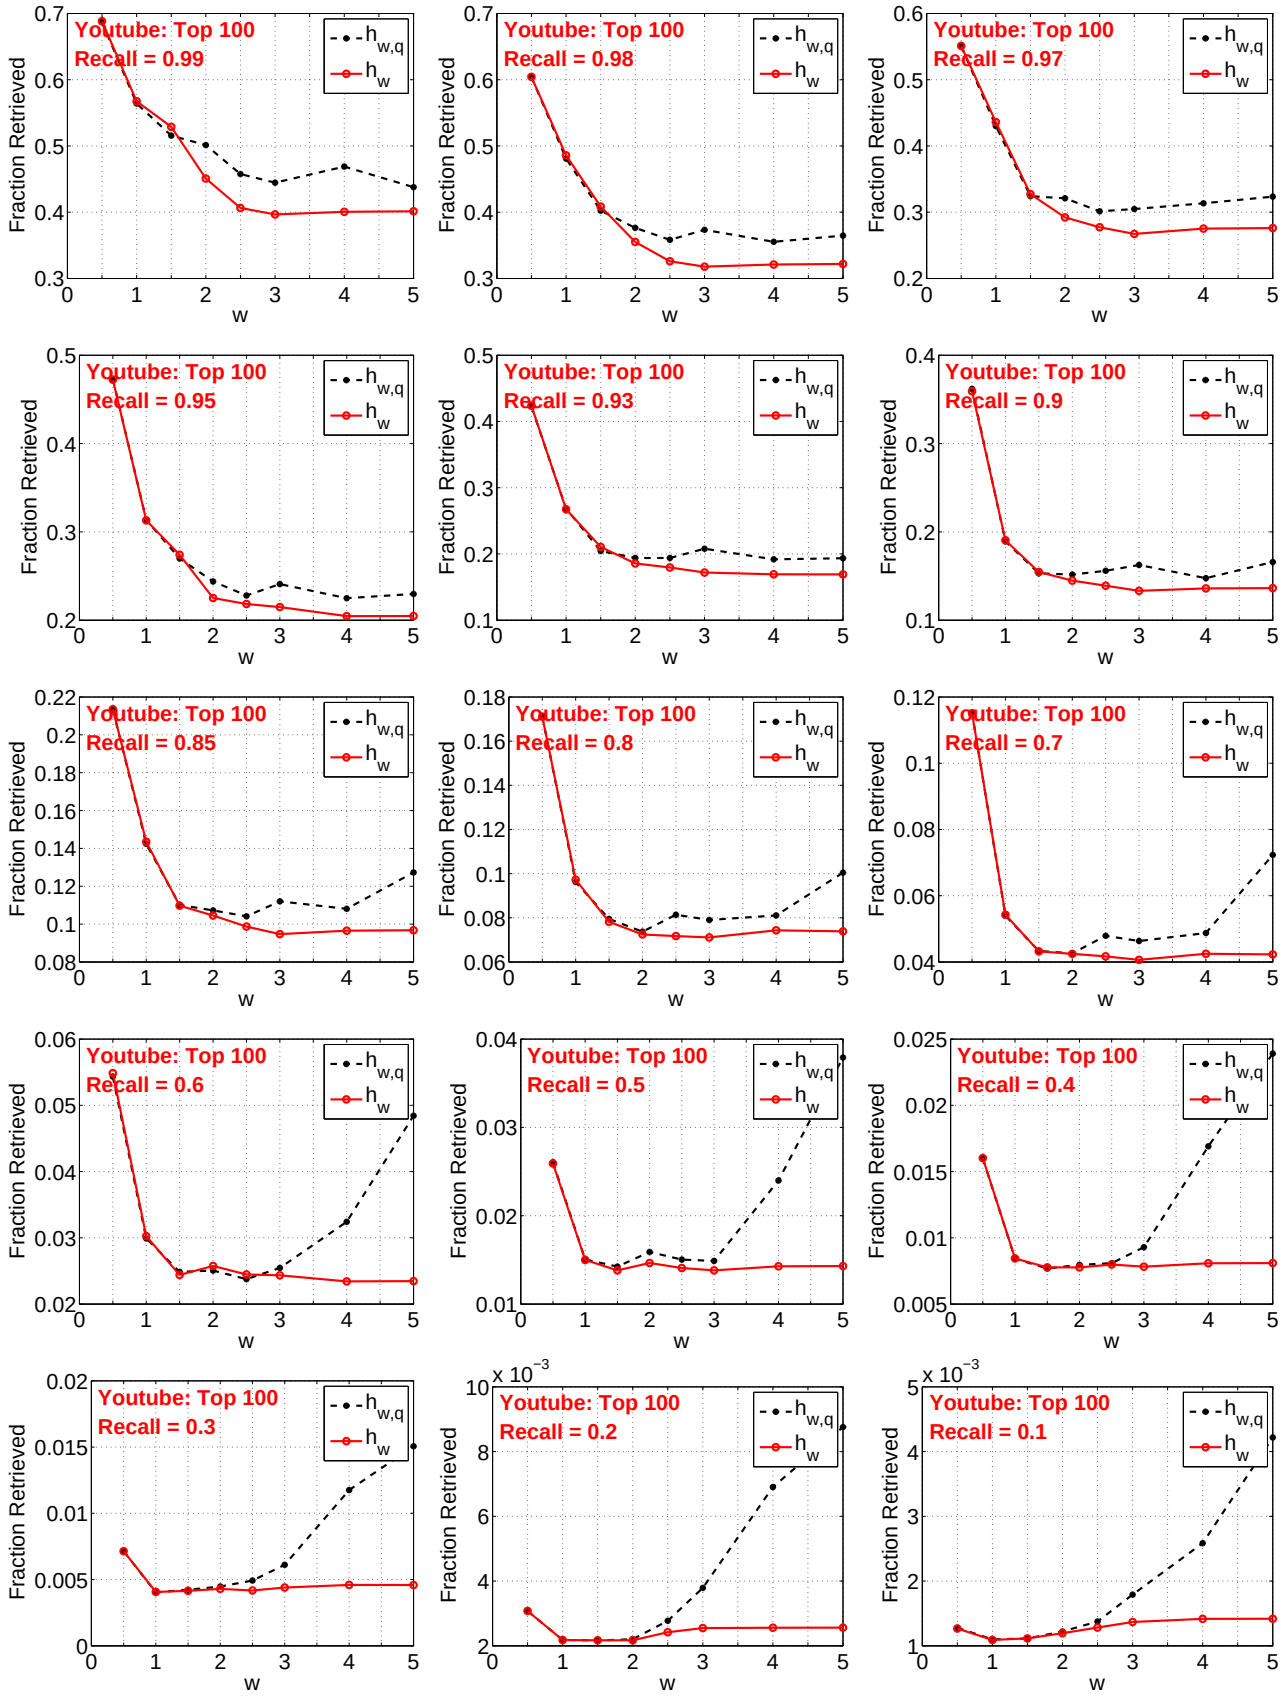

Figure 13: **Youtube Top 100** . In each panel, we plot the optimal *fraction retrieved* at a target *recall* value (for top-100) with respect to w for both coding schemes h_w and $h_{w,q}$. Lower is better.

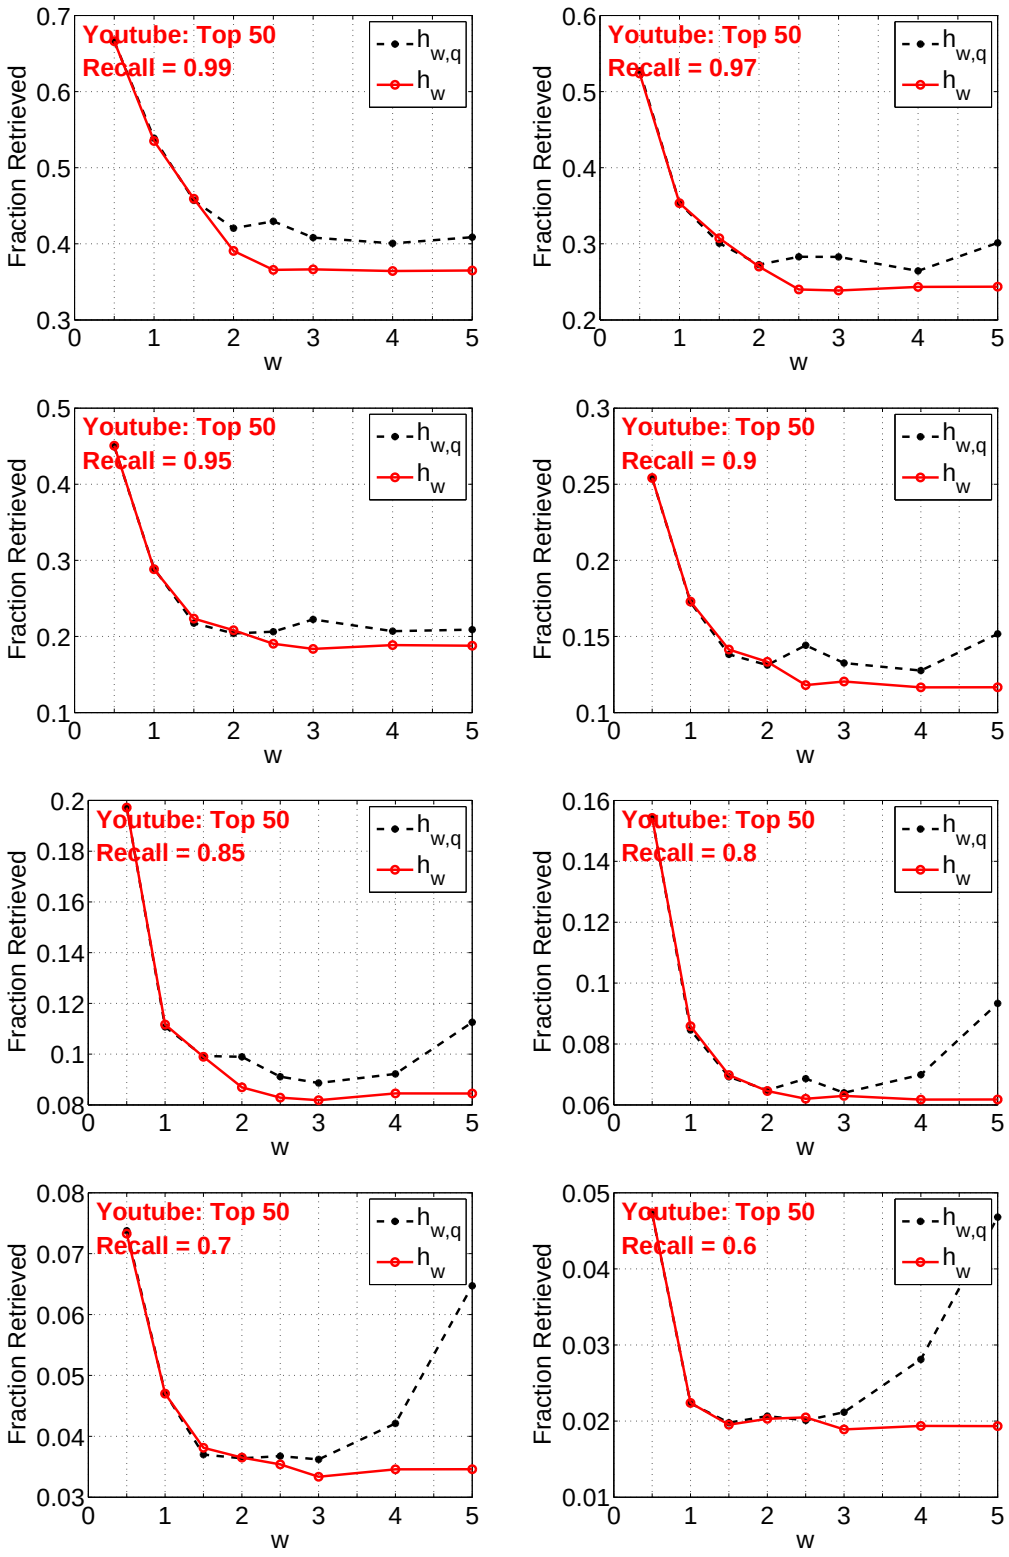

Figure 14: **Youtube Top 50** . In each panel, we plot the optimal *fraction retrieved* at a target *recall* value (for top-50) with respect to w for both coding schemes h_w and $h_{w,q}$.

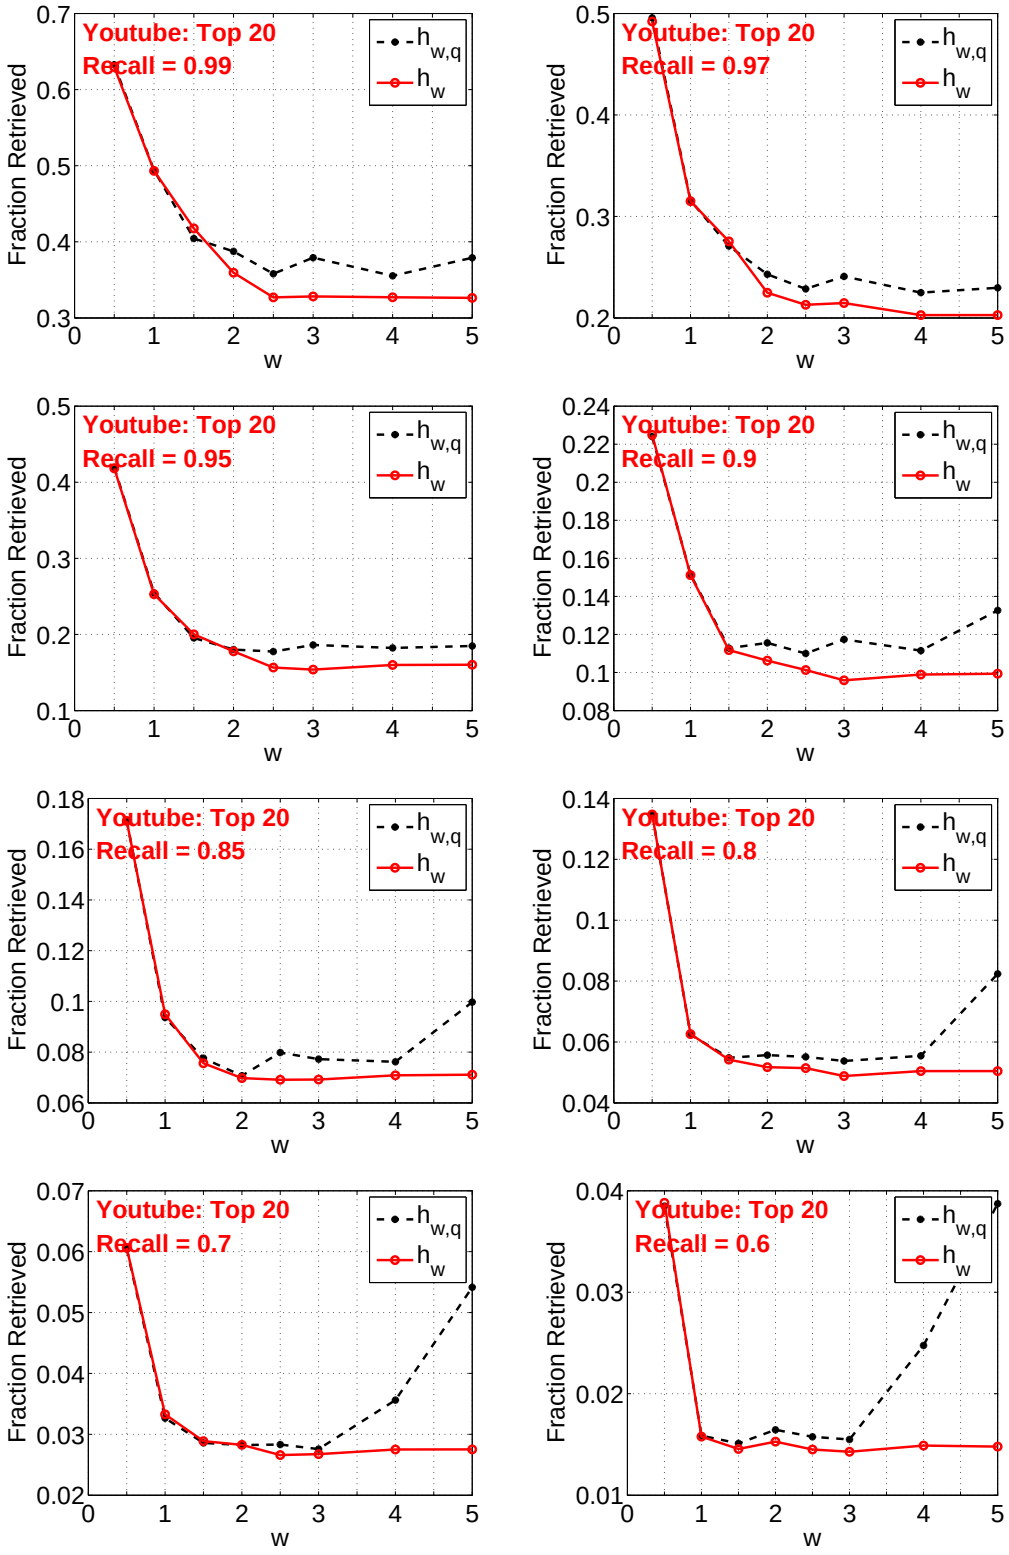

Figure 15: **Youtube Top 20** . In each panel, we plot the optimal *fraction retrieved* at a target *recall* value (for top-20) with respect to w for both coding schemes h_w and $h_{w,q}$.

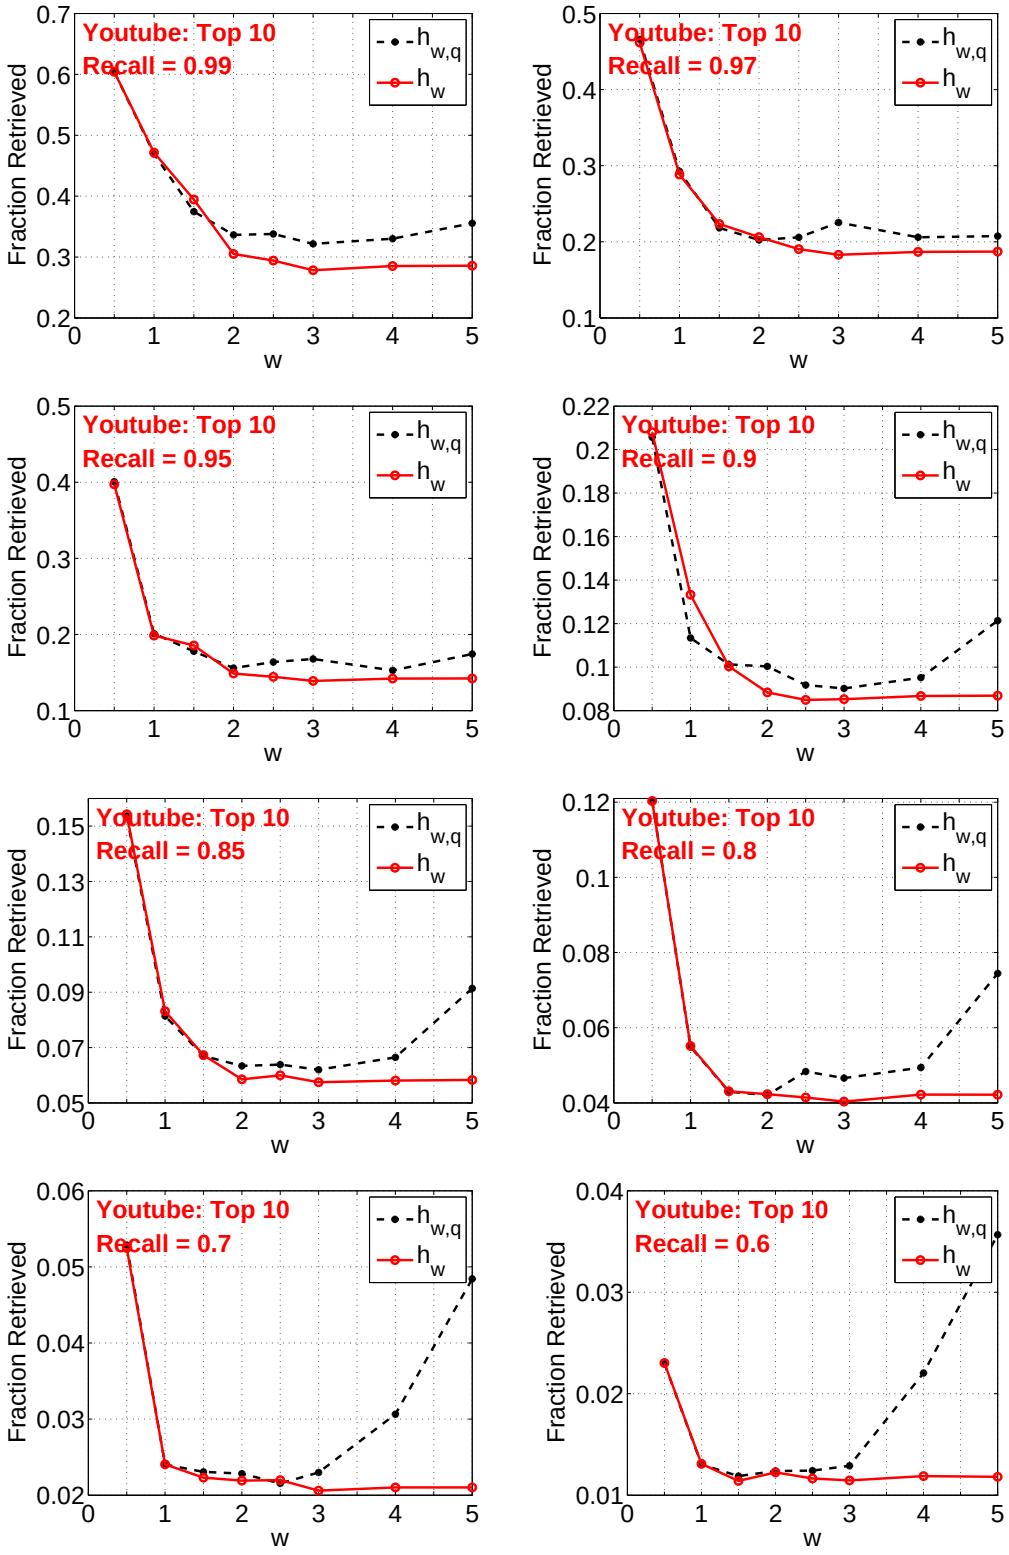

Figure 16: **Youtube Top 10** . In each panel, we plot the optimal *fraction retrieved* at a target *recall* value (for top-10) with respect to w for both coding schemes h_w and $h_{w,q}$.

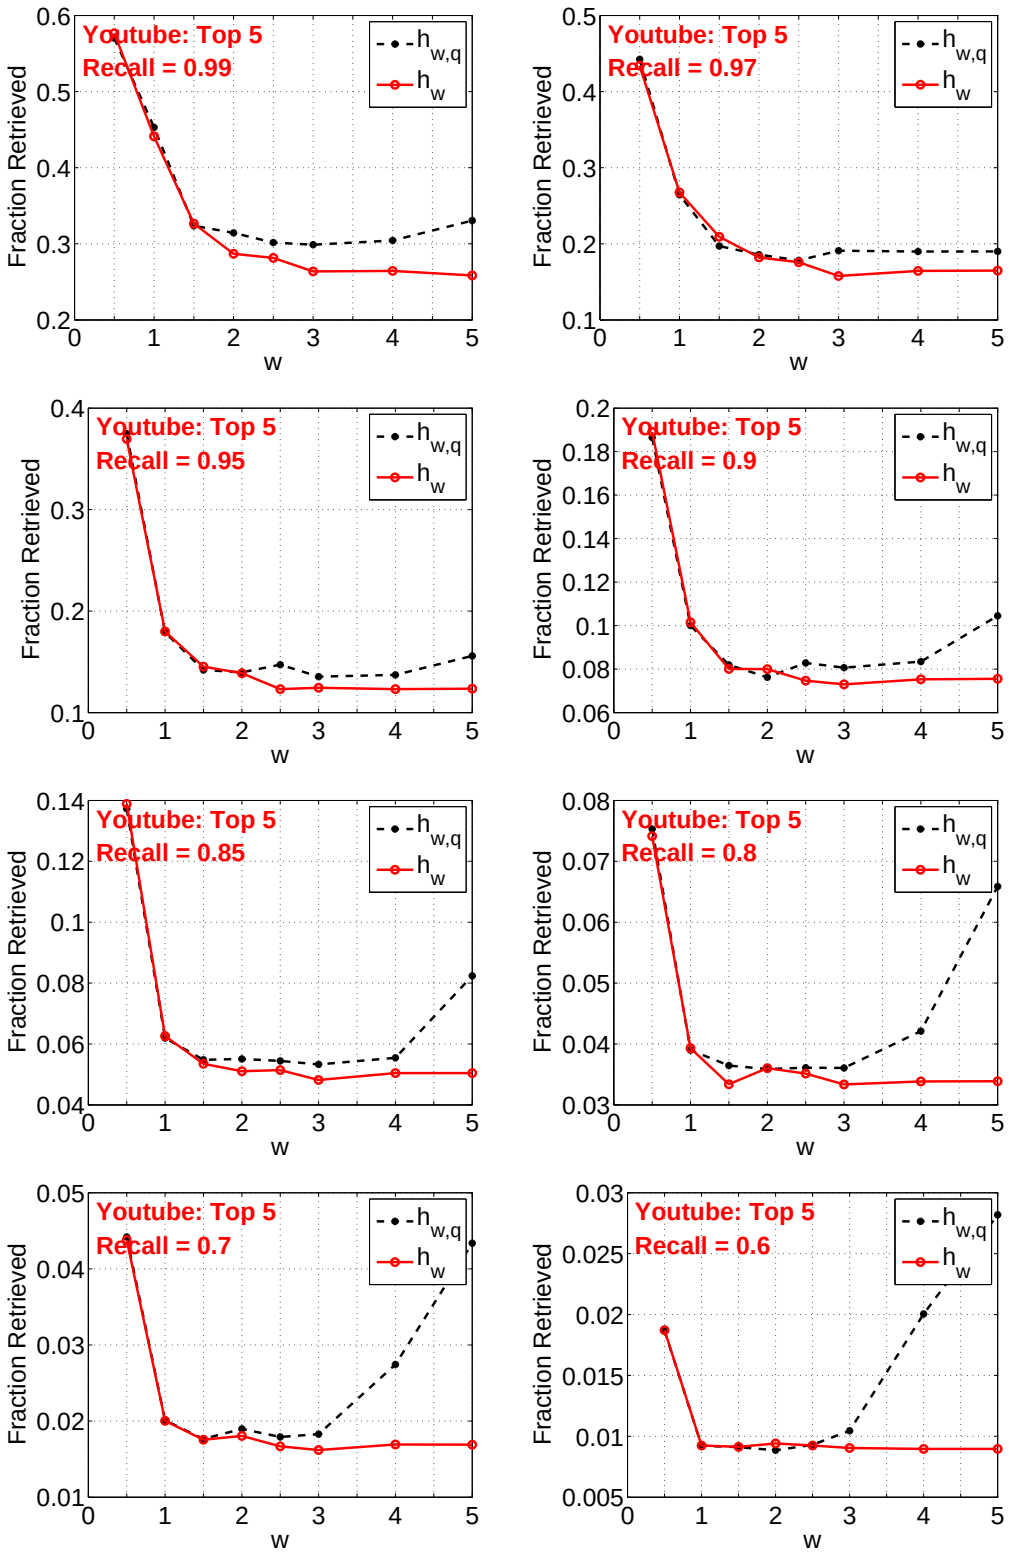

Figure 17: **Youtube Top 5** . In each panel, we plot the optimal *fraction retrieved* at a target *recall* value (for top-5) with respect to w for both coding schemes h_w and $h_{w,q}$.

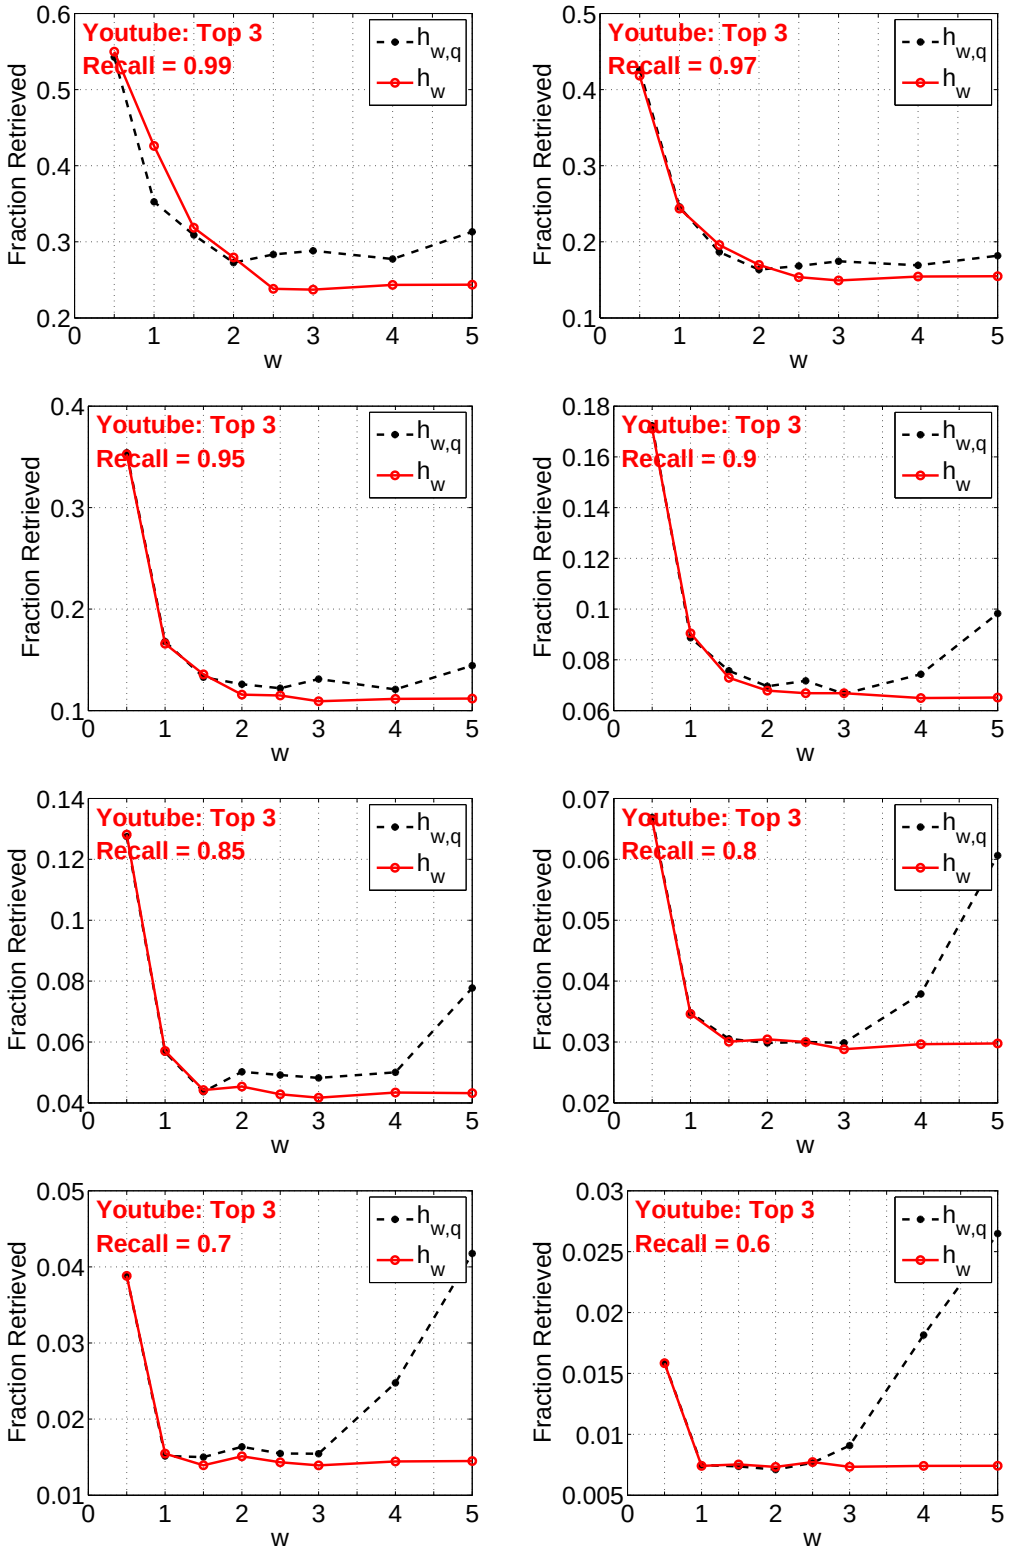

Figure 18: **Youtube Top 3** . In each panel, we plot the optimal *fraction retrieved* at a target *recall* value (for top-3) with respect to w for both coding schemes h_w and $h_{w,q}$.

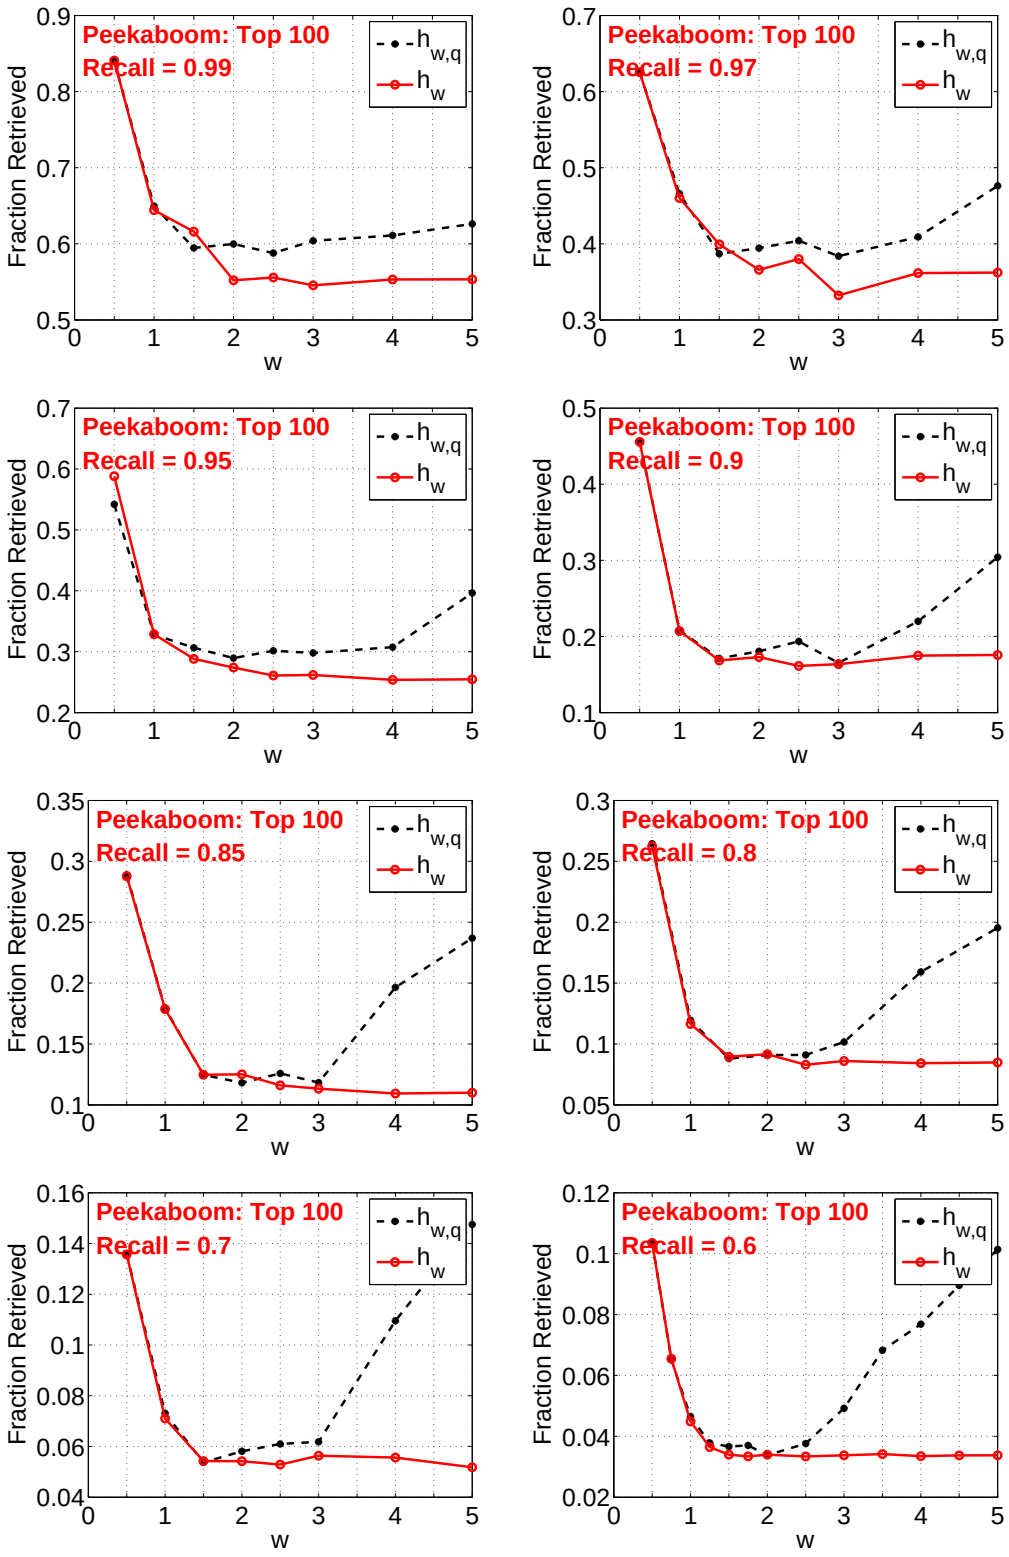

Figure 19: **Peekaboom Top 100** . In each panel, we plot the optimal *fraction retrieved* at a target *recall* value (for top-100) with respect to w for both coding schemes h_w and $h_{w,q}$.

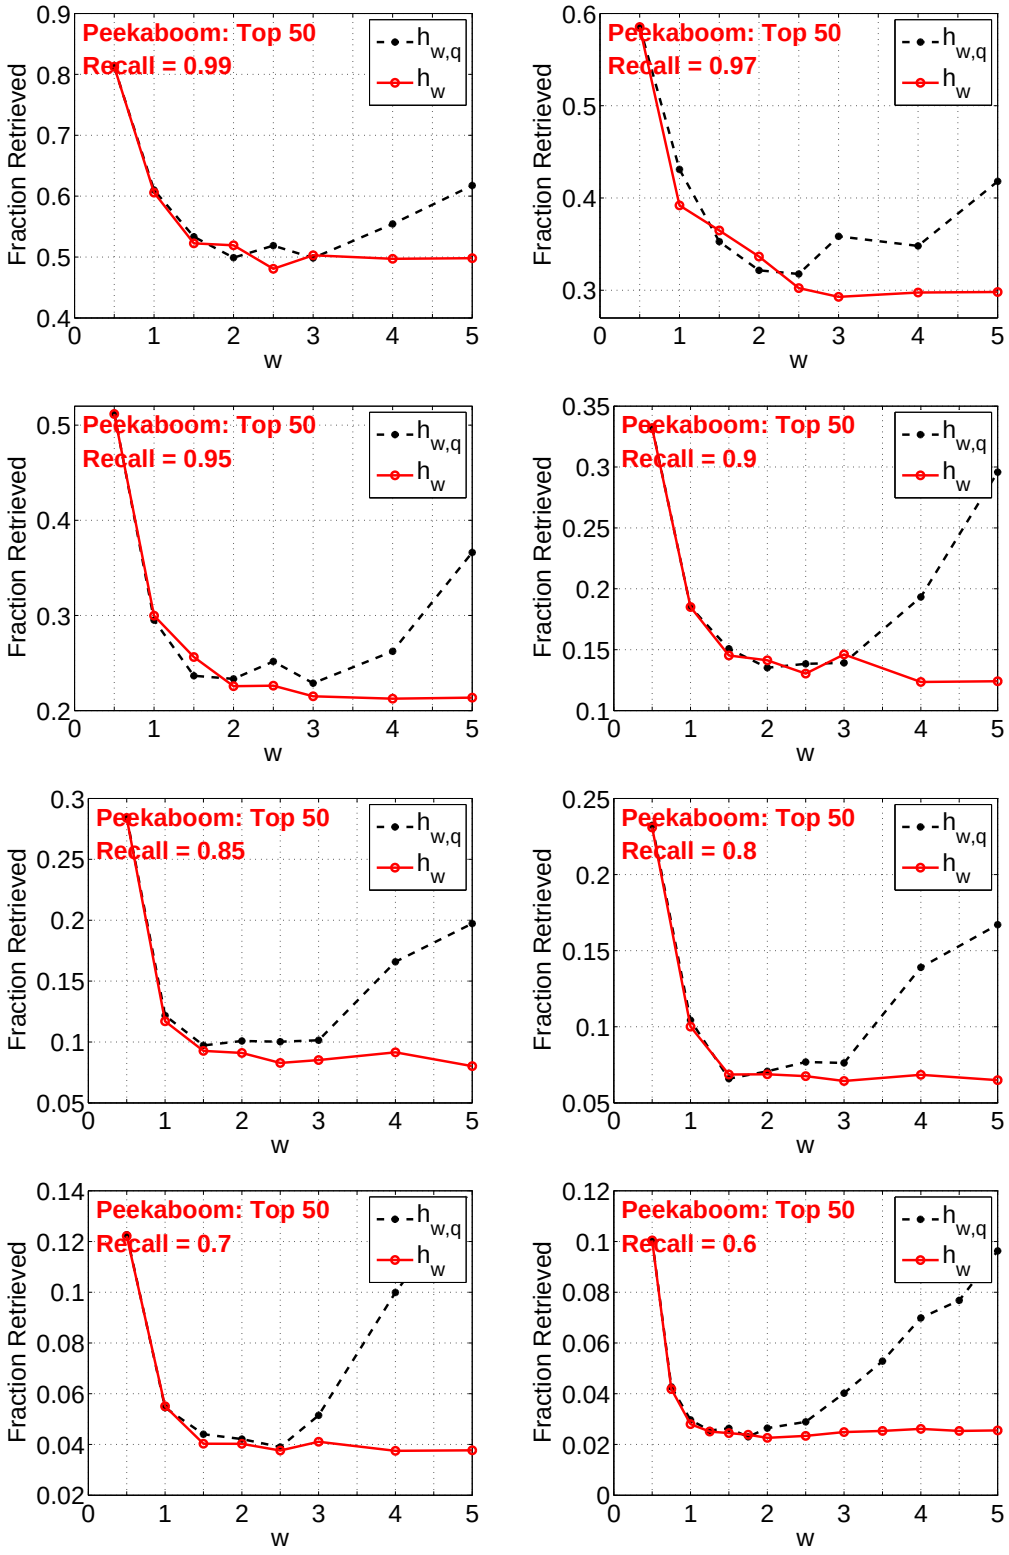

Figure 20: **Peekaboom Top 50** . In each panel, we plot the optimal *fraction retrieved* at a target *recall* value (for top-50) with respect to w for both coding schemes h_w and $h_{w,q}$.

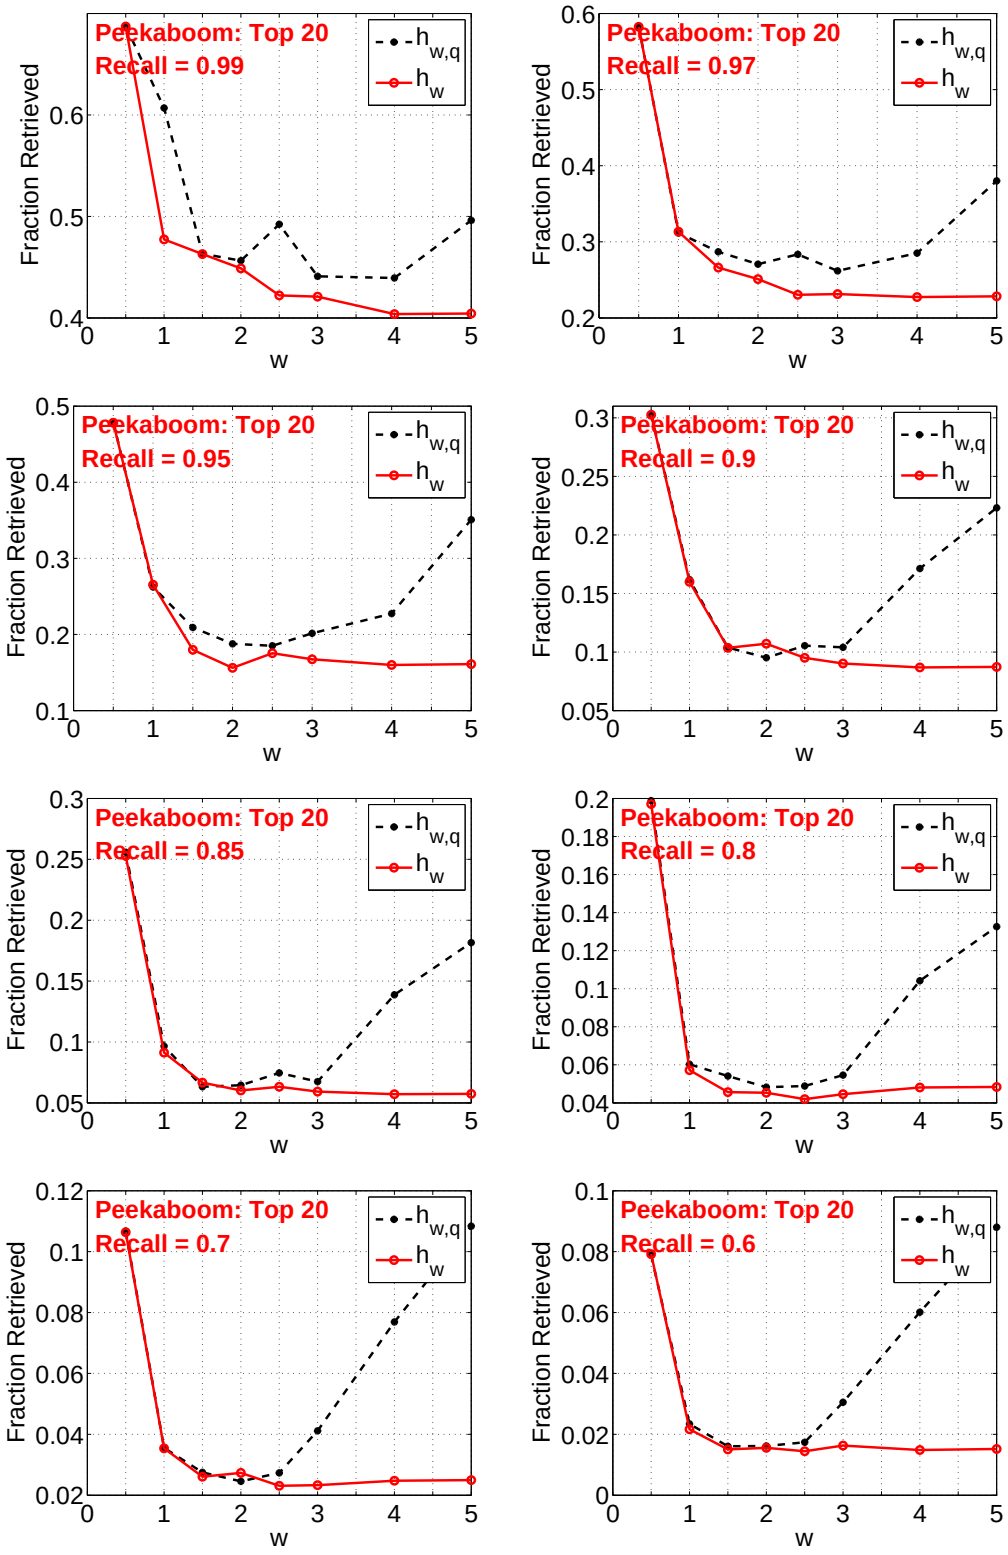

Figure 21: **Peekaboom Top 20** . In each panel, we plot the optimal *fraction retrieved* at a target *recall* value (for top-20) with respect to w for both coding schemes h_w and $h_{w,q}$.

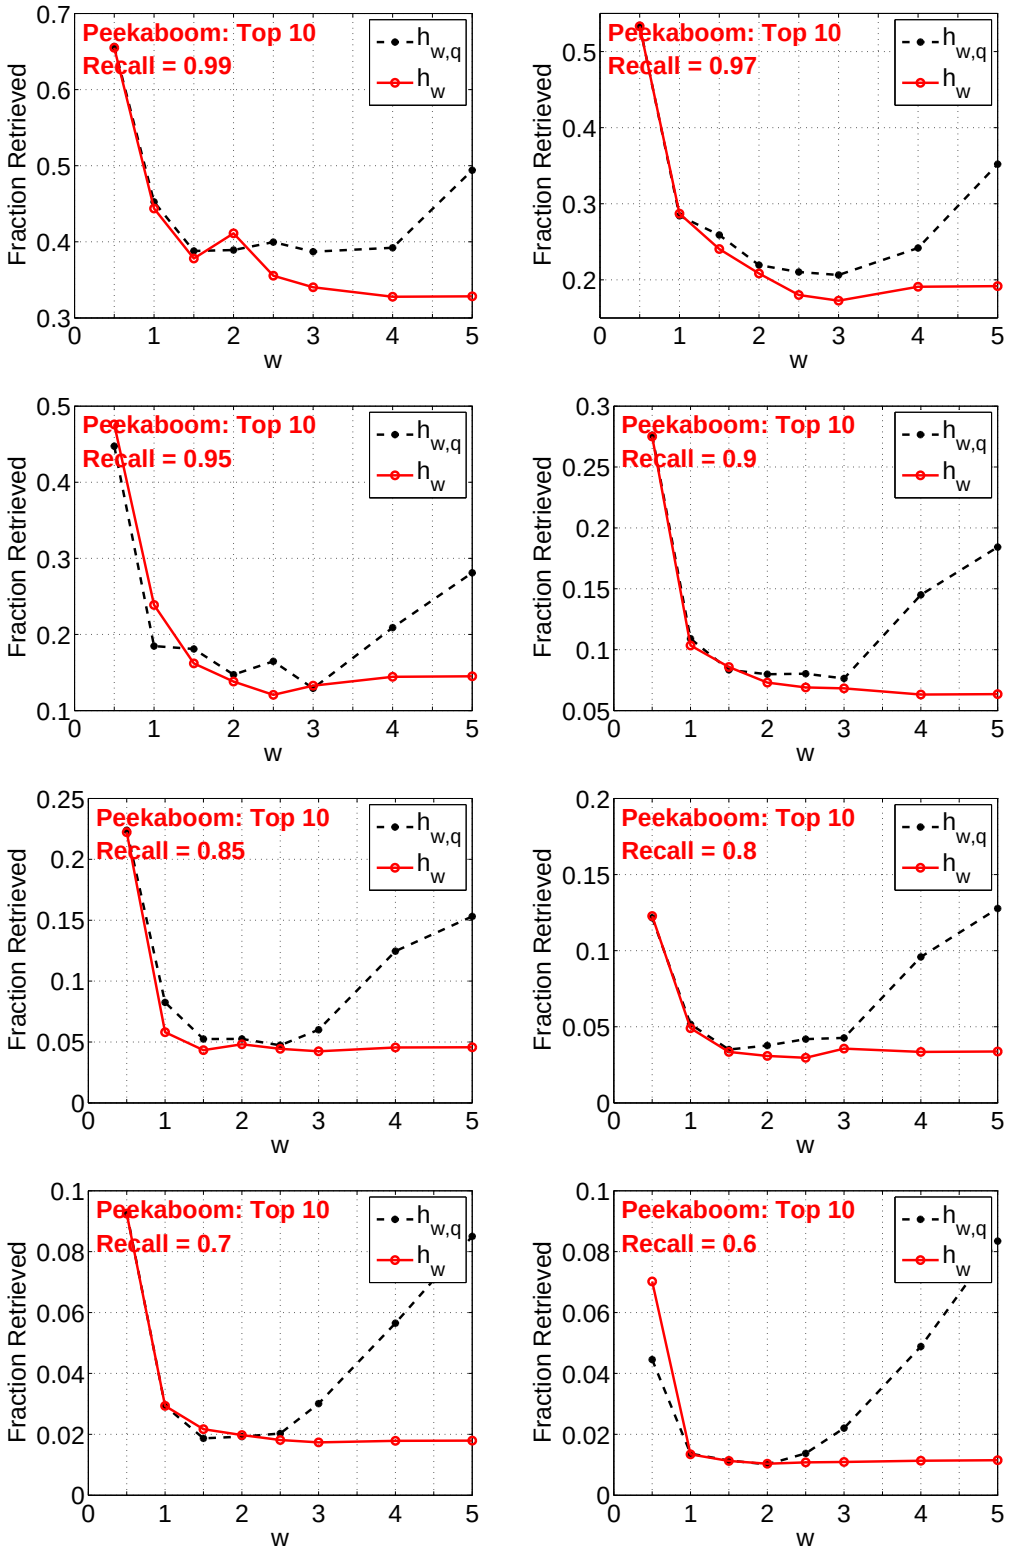

Figure 22: **Peekaboom Top 10** . In each panel, we plot the optimal *fraction retrieved* at a target *recall* value (for top-10) with respect to w for both coding schemes h_w and $h_{w,q}$.

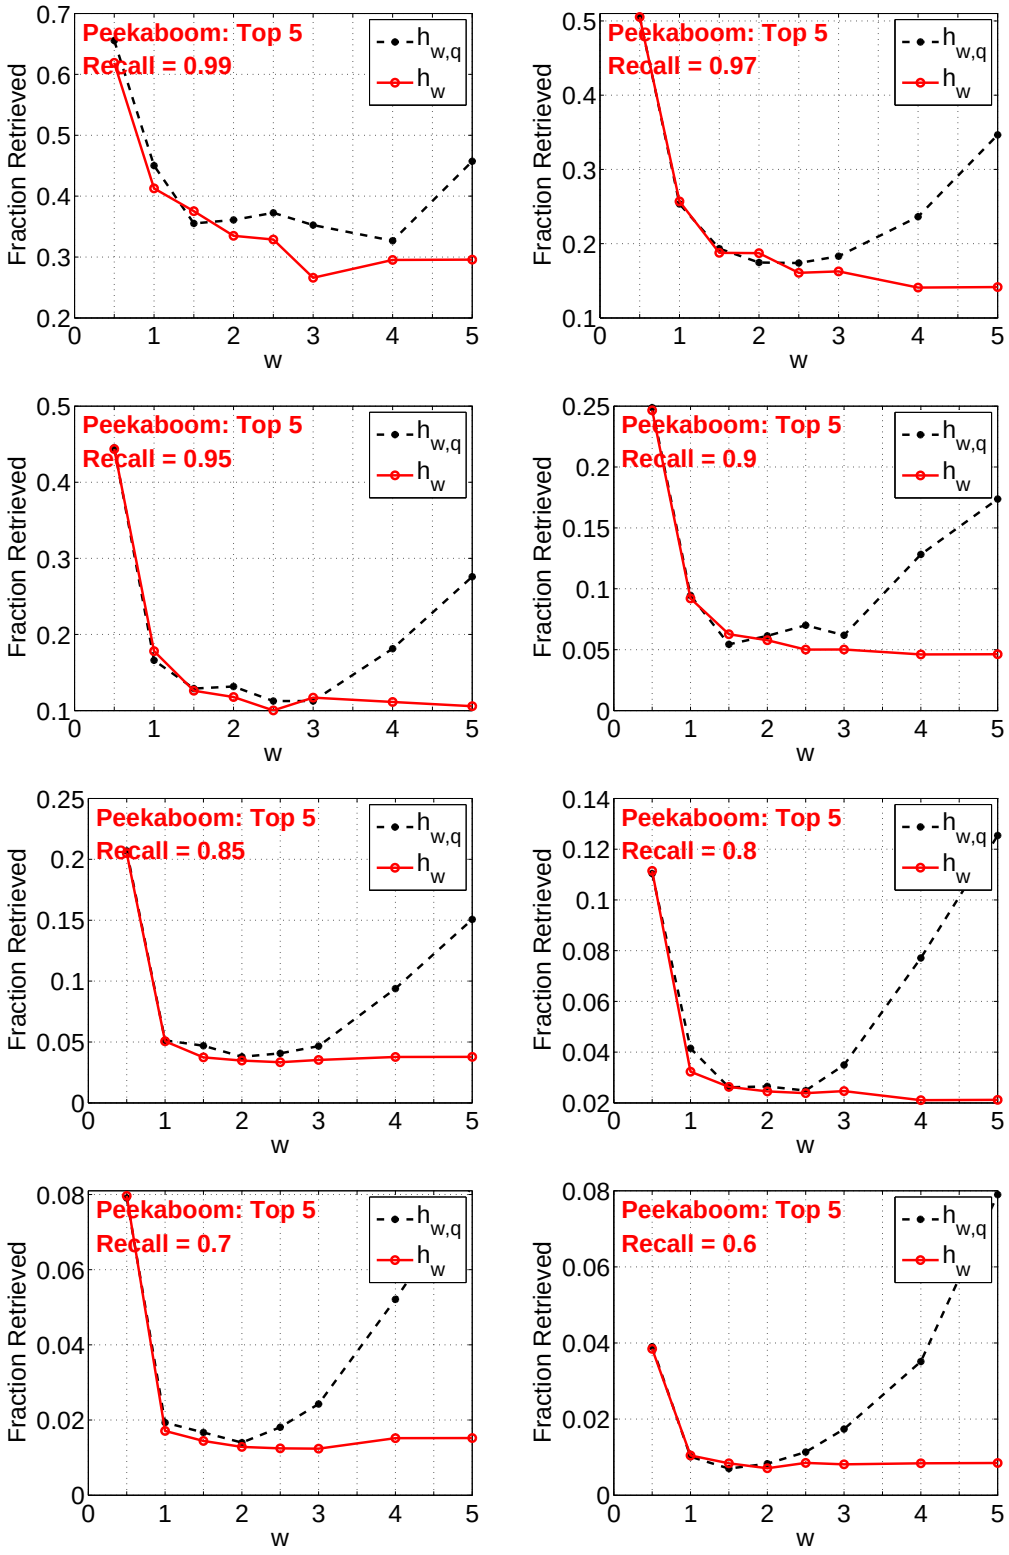

Figure 23: **Peekaboom Top 5** . In each panel, we plot the optimal *fraction retrieved* at a target *recall* value (for top-5) with respect to w for both coding schemes h_w and $h_{w,q}$.

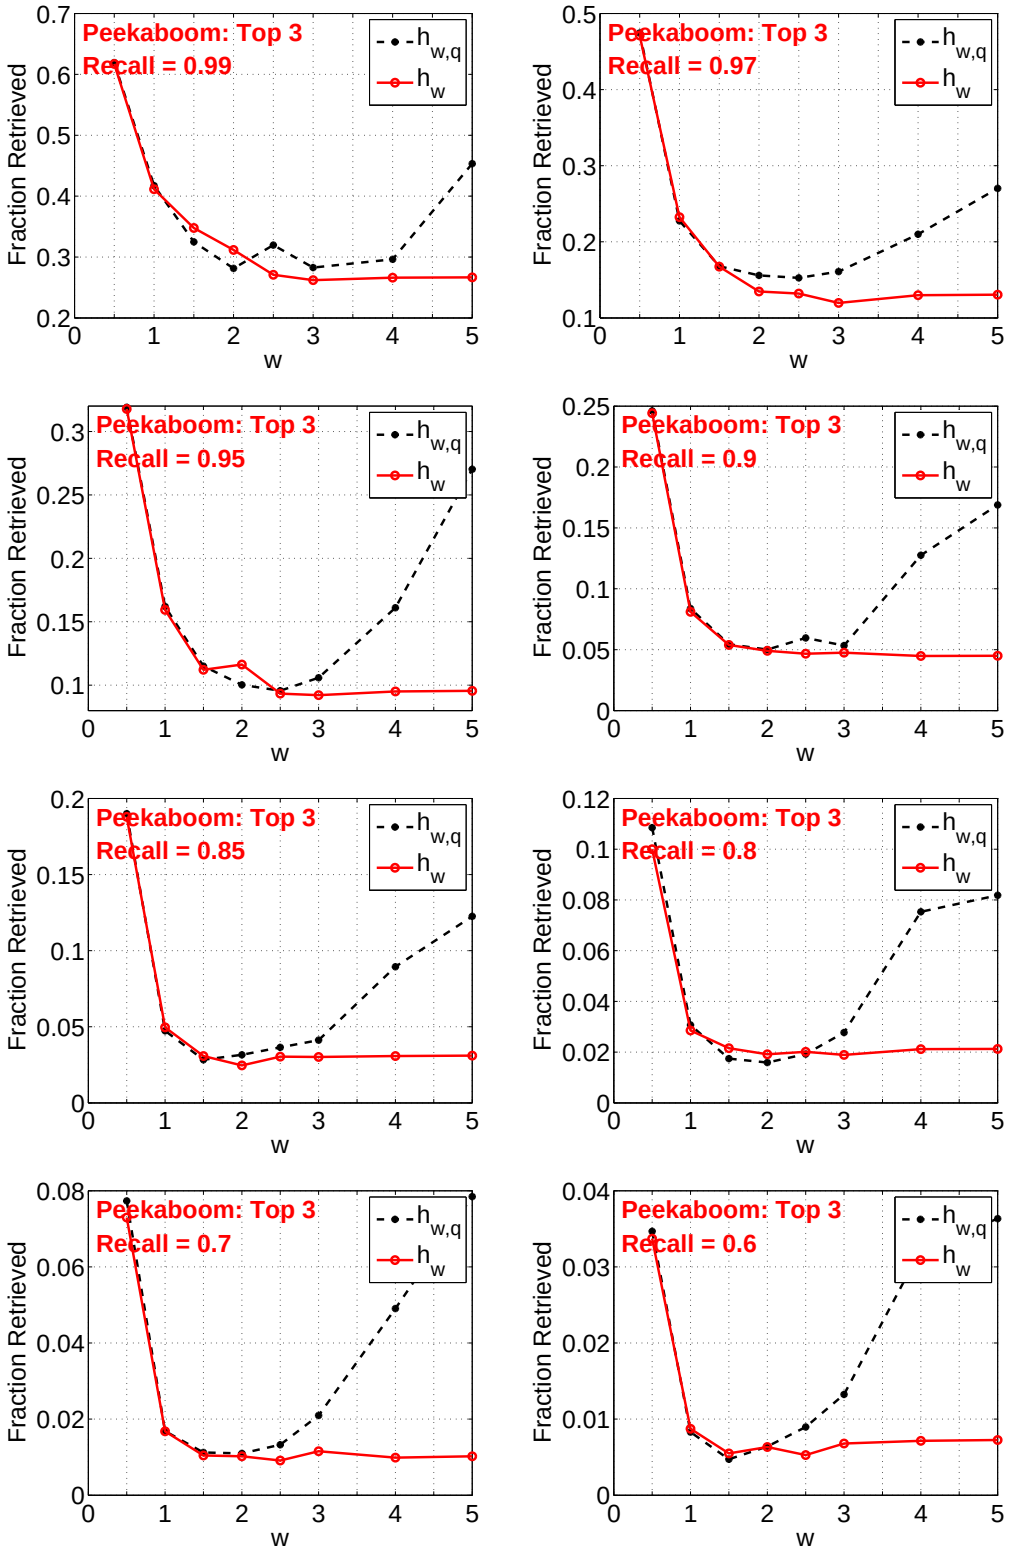

Figure 24: **Peekaboom Top 3** . In each panel, we plot the optimal *fraction retrieved* at a target *recall* value (for top-3) with respect to w for both coding schemes h_w and $h_{w,q}$.

6 Conclusion

We have compared two quantization (coding) schemes for random projections in the context of sublinear time approximate near neighbor search. The recently proposed scheme based on uniform quantization [4] is simpler than the influential existing work [1] (which used uniform quantization with a random offset). Our analysis confirms that, under the general theory of LSH, the new scheme [4] is simpler and more accurate than [1]. In other words, the step of random offset in [1] is not needed and may hurt the performance.

Our analysis provides the practical guidelines for using the proposed coding scheme to build hash tables. Our recommendation is to use a bin width about $w = 1.5$ when the target similarity is high and a bin width about $w = 3$ when the target similarity is not that high. In addition, using the proposed coding scheme based on uniform quantization (without the random offset), the influence of w is not very sensitive, which makes it very convenient in practical applications.

References

[1] Mayur Datar, Nicole Immorlica, Piotr Indyk, and Vahab S. Mirrokn. Locality-sensitive hashing scheme based on p-stable distributions. In *SCG*, pages 253 – 262, Brooklyn, NY, 2004.

[2] Jerome H. Friedman, F. Baskett, and L. Shustek. An algorithm for finding nearest neighbors. *IEEE Transactions on Computers*, 24:1000–1006, 1975.

[3] Piotr Indyk and Rajeev Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *STOC*, pages 604–613, Dallas, TX, 1998.

[4] Ping Li, Michael Mitzenmacher, and Anshumali Shrivastava. Coding for random projections. Technical report, arXiv:1308.2218, 2013.

[5] Ping Li, Michael Mitzenmacher, and Anshumali Shrivastava. Coding for random projections and non-linear estimators. Technical report, 2014.